



Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference, Release 5.1.x

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-30376-02



CONTENTS

27

CHAPTER 1

Border Gateway Protocol Commands 1

- accept-own 8
- additional-paths install backup 9
- additional-paths receive 11
- additional-paths selection 13
- additional-paths send 15
- address-family (BGP) 17
- advertise best-external 21
- advertise permanent-network 23
- advertisement-interval 24
- af-group 26
- aggregate-address 28
- aigp 31
- aigp send-cost-community 33
- allocate-label 35
- allowas-in 37
- as-format 39
- as-override 40
- as-path-loopcheck out disable 42
- attribute-filter group 44
- bfd (BGP) 46
- bgp as-path-loopcheck 51
- bgp attribute-download 52
- bgp auto-policy-soft-reset disable 54
- bgp bestpath as-path ignore 56
- bgp bestpath compare-routerid 58
- bgp bestpath cost-community ignore 60

| | |
|--|-----|
| bgp bestpath med always | 62 |
| bgp bestpath med confed | 64 |
| bgp bestpath med missing-as-worst | 66 |
| bgp client-to-client reflection disable | 68 |
| bgp cluster-id | 70 |
| bgp confederation identifier | 72 |
| bgp confederation peers | 74 |
| bgp dampening | 76 |
| bgp default local-preference | 78 |
| bgp enforce-first-as disable | 79 |
| bgp fast-external-fallover disable | 81 |
| bgp graceful-restart | 82 |
| bgp graceful-restart graceful-reset | 84 |
| bgp graceful-restart purge-time | 86 |
| bgp graceful-restart restart-time | 88 |
| bgp graceful-restart stalepath-time | 90 |
| bgp import-delay | 92 |
| bgp label-delay | 94 |
| bgp log neighbor changes disable | 96 |
| bgp maximum neighbor | 98 |
| bgp redistribute-internal | 99 |
| bgp router-id | 101 |
| bgp scan-time | 103 |
| bgp update-delay | 105 |
| bgp write-limit | 107 |
| capability additional-paths receive | 109 |
| capability additional-paths send | 111 |
| capability orf prefix | 113 |
| capability suppress 4-byte-as | 116 |
| clear bgp | 119 |
| clear bgp dampening | 121 |
| clear bgp external | 123 |
| clear bgp flap-statistics | 125 |
| clear bgp nexthop performance-statistics | 127 |
| clear bgp nexthop registration | 129 |

[clear bgp peer-drops](#) 131

[clear bgp performance-statistics](#) 132

[clear bgp self-originated](#) 133

[clear bgp shutdown](#) 135

[clear bgp soft](#) 137

[default-information originate \(BGP\)](#) 140

[default-martian-check disable](#) 142

[default-metric \(BGP\)](#) 143

[default-originate](#) 144

[description \(BGP\)](#) 146

[distance bgp](#) 148

[distribute bgp-ls \(ISIS\)](#) 150

[distribute bgp-ls \(OSPF\)](#) 152

[domain-distinguisher](#) 153

[dmz-link-bandwidth](#) 154

[dscp \(BGP\)](#) 156

[ebgp-multihop](#) 159

[export route-policy](#) 161

[export route-target](#) 163

[import route-policy](#) 165

[import route-target](#) 167

[ignore-connected-check](#) 169

[keychain](#) 170

[keychain-disable](#) 172

[keychain inheritance-disable](#) 174

[label-allocation-mode](#) 176

[label mode](#) 178

[local-as](#) 181

[maximum-paths \(BGP\)](#) 183

[maximum-prefix \(BGP\)](#) 185

[mpls activate \(BGP\)](#) 188

[mvpn](#) 191

[neighbor \(BGP\)](#) 192

[neighbor-group](#) 194

[network \(BGP\)](#) 196

- network backdoor 198
- next-hop-self 200
- next-hop-unchanged 203
- nexthop resolution prefix-length minimum 205
- nexthop route-policy 206
- nexthop trigger-delay 208
- nsr (BGP) 210
- orf 212
- password (BGP) 214
- password (rpki-server) 216
- password-disable 218
- permanent-network 220
- precedence 222
- preference (rpki-server) 224
- purge-time (rpki-server) 225
- rd 227
- receive-buffer-size 229
- redistribute (BGP) 231
- refresh-time (rpki-server) 235
- response-time (rpki-server) 237
- remote-as (BGP) 239
- remove-private-as 242
- retain local-label 245
- retain route-target 247
- route-policy (BGP) 249
- route-reflector-client 251
- router bgp 253
- rpki server 255
- selective-vrf-download disable 256
- send-buffer-size 258
- send-community-ebgp 260
- send-extended-community-ebgp 262
- session-group 264
- session-open-mode 266
- show bgp 268

show bgp update out 282

show bgp update in error process 284

show bgp update out filter-group 285

show bgp update out process 286

show bgp update out sub-group 288

show bgp update out update-group 290

show bgp vrf update in error 292

show bgp advertised 293

show bgp af-group 301

show bgp attribute-key 304

show bgp cidr-only 309

show bgp community 313

show bgp convergence 318

show bgp dampened-paths 321

show bgp flap-statistics 326

show bgp inconsistent-as 332

show bgp labels 337

show bgp l2vpn rd 341

show bgp l2vpn vpls 344

show bgp neighbor-group 349

show bgp neighbors 353

show bgp neighbors nsr 375

show bgp nexthops 377

show bgp nsr 386

show bgp paths 390

show bgp policy 393

show bgp process 402

show bgp regexp 426

show bgp route-policy 431

show bgp session-group 436

show bgp sessions 439

show bgp summary 442

show bgp summary nsr 447

show bgp truncated-communities 451

show bgp update-group 456

show bgp vrf imported-routes 464
show protocols (BGP) 467
show svd role 470
show svd state 472
shutdown (BGP) 473
shutdown (rpki-server) 475
signalling disable 476
site-of-origin (BGP) 477
socket receive-buffer-size 479
socket send-buffer-size 481
soft-reconfiguration inbound 483
speaker-id 486
svd platform enable 487
table-policy 489
timers (BGP) 491
timers bgp 493
transport (rpki-server) 495
ttl-security 497
update limit 500
update limit address-family 502
update limit sub-group 504
update in error-handling basic disable 506
update in error-handling extended 508
update out logging 510
update-source 511
update wait-install 513
use 514
username (rpki-server) 519
vrf (BGP) 520
weight 521

CHAPTER 2**BFD Commands 525**

address-family ipv4 unicast (BFD) 527
bfd 529
bfd address-family ipv4 destination 531

bfd address-family ipv4 fast-detect 533
bfd address-family ipv4 minimum-interval 534
bfd address-family ipv4 multiplier 537
bfd address-family ipv4 timers 540
bundle coexistence bob-blb 542
bfd dampening 544
bfd fast-detect 546
bfd minimum-interval 549
bfd multipath include location 553
bfd multiplier 555
clear bfd counters 558
dampening (BFD) 560
echo disable 563
echo ipv4 source 565
echo latency detect 567
echo startup validate 569
interface (BFD) 571
ipv6 checksum 574
multihop ttl-drop-threshold 576
show bfd 578
show bfd client 581
show bfd counters 583
show bfd mib session 586
show bfd multipath 589
show bfd neighbor 591
show bfd session 593
show bfd summary 601

CHAPTER 3**Enhanced Interior Gateway Routing Protocol Commands 603**

address-family (EIGRP) 605
authentication keychain 607
auto-summary (EIGRP) 609
autonomous-system 611
bandwidth-percent (EIGRP) 613
clear eigrp neighbors 615

| | |
|-------------------------------|-----|
| clear eigrp topology | 617 |
| default-information | 619 |
| default-metric (EIGRP) | 621 |
| distance (EIGRP) | 623 |
| hello-interval (EIGRP) | 625 |
| hold-time (EIGRP) | 626 |
| interface (EIGRP) | 628 |
| log-neighbor-changes | 630 |
| log-neighbor-warnings | 631 |
| maximum-paths (EIGRP) | 632 |
| maximum-prefix (EIGRP) | 634 |
| metric (EIGRP) | 637 |
| metric maximum-hops | 639 |
| metric rib-scale | 641 |
| metric weights | 643 |
| neighbor (EIGRP) | 646 |
| neighbor maximum-prefix | 648 |
| next-hop-self disable | 651 |
| nsf disable (EIGRP) | 652 |
| passive-interface (EIGRP) | 654 |
| redistribute (EIGRP) | 655 |
| redistribute maximum-prefix | 658 |
| route-policy (EIGRP) | 660 |
| router eigrp | 662 |
| router-id (EIGRP) | 664 |
| show eigrp accounting | 665 |
| show eigrp interfaces | 668 |
| show eigrp neighbors | 673 |
| show eigrp topology | 677 |
| show eigrp traffic | 681 |
| show protocols (EIGRP) | 683 |
| site-of-origin (EIGRP) | 687 |
| split-horizon disable (EIGRP) | 689 |
| stub (EIGRP) | 690 |
| summary-address (EIGRP) | 693 |

timers active-time **695**
timers nsf route-hold (EIGRP) **697**
variance **699**
vrf (EIGRP) **701**

CHAPTER 4**IS-IS Commands 703**

address-family (IS-IS) **707**
address-family multicast topology (IS-IS) **709**
adjacency-check disable **711**
circuit-type **713**
clear isis process **715**
clear isis route **717**
clear isis statistics **719**
csnp-interval **721**
default-information originate (IS-IS) **723**
disable (IS-IS) **725**
distance (IS-IS) **726**
fast-reroute per-link (IS-IS) **728**
fast-reroute per-prefix (IS-IS) **730**
fast-reroute per-link priority-limit (IS-IS) **732**
fast-reroute per-prefix load-sharing disable (IS-IS) **734**
fast-reroute per-prefix tiebreaker (IS-IS) **735**
hello-interval (IS-IS) **737**
hello-multiplier **739**
hello-padding **741**
hello-password **743**
hello-password keychain **745**
hello-password accept **747**
hostname dynamic disable **749**
ignore-lsp-errors **751**
interface (IS-IS) **753**
ispf **755**
is-type **756**
log adjacency changes (IS-IS) **758**
log pdu drops **760**

| | |
|---|-----|
| lsp fast-flood threshold | 761 |
| lsp-gen-interval | 763 |
| lsp-interval | 765 |
| lsp-mtu | 767 |
| lsp-password | 769 |
| lsp-password accept | 771 |
| lsp-refresh-interval | 773 |
| maximum-paths (IS-IS) | 775 |
| maximum-redistributed-prefixes (IS-IS) | 776 |
| max-lsp-lifetime | 778 |
| mesh-group (IS-IS) | 780 |
| metric (IS-IS) | 782 |
| metric-style narrow | 784 |
| metric-style transition | 786 |
| metric-style wide | 788 |
| min-lsp-arrivaltime | 790 |
| mpls ldp auto-config | 792 |
| mpls ldp sync (IS-IS) | 794 |
| mpls traffic-eng (IS-IS) | 796 |
| mpls traffic-eng multicast-intact (IS-IS) | 798 |
| mpls traffic-eng path-selection ignore overload | 800 |
| mpls traffic-eng router-id (IS-IS) | 802 |
| net | 804 |
| nsf (IS-IS) | 806 |
| nsf interface-expires | 808 |
| nsf interface-timer | 810 |
| nsf lifetime (IS-IS) | 812 |
| passive (IS-IS) | 813 |
| point-to-point | 814 |
| priority (IS-IS) | 815 |
| propagate level | 817 |
| redistribute (IS-IS) | 819 |
| retransmit-interval (IS-IS) | 823 |
| retransmit-throttle-interval | 825 |
| router isis | 827 |

set-attached-bit 829
set-overload-bit 831
show isis 833
show isis adjacency 836
show isis adjacency-log 839
show isis checkpoint adjacency 841
show isis checkpoint interface 843
show isis checkpoint lsp 845
show isis database 847
show isis database-log 856
show isis fast-reroute 859
show isis hostname 861
show isis interface 863
show isis lsp-log 868
show isis mesh-group 871
show isis mpls traffic-eng adjacency-log 873
show isis mpls traffic-eng advertisements 875
show isis mpls traffic-eng tunnel 878
show isis neighbors 880
show isis protocol 884
show isis route 887
show isis spf-log 890
show isis statistics 898
show isis topology 902
show protocols (IS-IS) 905
shutdown (IS-IS) 909
single-topology 910
snmp-server traps isis 912
spf-interval 913
spf prefix-priority (IS-IS) 915
summary-prefix (IS-IS) 917
suppressed 919
tag (IS-IS) 921
topology-id 923
trace (IS-IS) 925

CHAPTER 5**OSPF Commands 927**

- address-family (OSPF) 931
- adjacency stagger 932
- area (OSPF) 934
- authentication (OSPF) 936
- authentication-key (OSPF) 939
- auto-cost (OSPF) 941
- capability opaque disable 943
- capability vrf-lite (OSPFv2) 944
- clear ospf process 946
- clear ospf redistribution 948
- clear ospf routes 950
- clear ospf statistics 952
- cost (OSPF) 954
- cost-fallback (OSPF) 956
- database-filter all out (OSPF) 958
- dead-interval (OSPF) 960
- default-cost (OSPF) 962
- default-information originate (OSPF) 964
- default-metric (OSPF) 966
- demand-circuit (OSPF) 968
- disable-dn-bit-check 970
- distance (OSPF) 971
- distance ospf 974
- distribute-list 976
- domain-id (OSPF) 978
- domain-tag 980
- fast-reroute (OSPFv2) 981
- fast-reroute per-link exclude interface 983
- fast-reroute per-prefix exclude interface (OSPFv2) 985
- fast-reroute per-prefix lfa-candidate (OSPFv2) 987
- fast-reroute per-prefix remote-lfa (OSPFv2) 988
- fast-reroute per-prefix use-candidate-only (OSPFv2) 990
- flood-reduction (OSPF) 991

hello-interval (OSPF) 993

ignore lsa mospf 995

interface (OSPF) 996

log adjacency changes (OSPF) 998

loopback stub-network 999

max-lsa 1001

max-metric 1004

maximum interfaces (OSPF) 1007

maximum paths (OSPF) 1009

maximum redistributed-prefixes (OSPF) 1011

message-digest-key 1013

mpls ldp auto-config (OSPF) 1016

mpls ldp sync (OSPF) 1017

mpls traffic-eng (OSPF) 1018

mpls traffic-eng igp-intact (OSPF) 1020

mpls traffic-eng multicast-intact (OSPF) 1022

mpls traffic-eng router-id (OSPF) 1024

mtu-ignore (OSPF) 1026

multi-area-interface 1028

neighbor (OSPF) 1030

neighbor database-filter all out 1033

network (OSPF) 1035

nsf (OSPF) 1037

nsf flush-delay-time (OSPF) 1039

nsf interval (OSPF) 1040

nsf lifetime (OSPF) 1042

nsr (OSPF) 1044

nssa (OSPF) 1045

ospf name-lookup 1047

packet-size (OSPF) 1048

passive (OSPF) 1050

priority (OSPF) 1052

protocol shutdown 1054

queue dispatch flush-lsa 1055

queue dispatch incoming 1057

| | |
|---------------------------------|------|
| queue dispatch rate-limited-lsa | 1059 |
| queue dispatch spf-lsa-limit | 1061 |
| queue limit | 1063 |
| range (OSPF) | 1065 |
| redistribute (OSPF) | 1067 |
| retransmit-interval (OSPF) | 1072 |
| route-policy (OSPF) | 1074 |
| router-id (OSPF) | 1076 |
| router ospf | 1078 |
| security ttl (OSPF) | 1080 |
| sham-link | 1082 |
| show ospf | 1084 |
| show ospf border-routers | 1088 |
| show ospf database | 1090 |
| show ospf flood-list | 1104 |
| show ospf interface | 1106 |
| show ospf mpls traffic-eng | 1109 |
| show ospf message-queue | 1114 |
| show ospf neighbor | 1117 |
| show ospf request-list | 1125 |
| show ospf retransmission-list | 1128 |
| show ospf routes | 1131 |
| show ospf sham-links | 1136 |
| show ospf summary-prefix | 1139 |
| show ospf virtual-links | 1141 |
| show protocols (OSPF) | 1143 |
| snmp context (OSPF) | 1146 |
| snmp trap (OSPF) | 1149 |
| snmp trap rate-limit (OSPF) | 1150 |
| spf prefix-priority (OSPFv2) | 1151 |
| stub (OSPF) | 1153 |
| summary-prefix (OSPF) | 1155 |
| timers lsa group-pacing | 1157 |
| timers lsa min-arrival | 1159 |
| timers throttle lsa all (OSPF) | 1160 |

timers throttle spf (OSPF) 1163
transmit-delay (OSPF) 1165
virtual-link (OSPF) 1167
vrf (OSPF) 1169

CHAPTER 6**OSPFv3 Commands 1171**

address-family (OSPFv3) 1174
area (OSPFv3) 1175
authentication (OSPFv3) 1177
auto-cost (OSPFv3) 1179
capability vrf-lite (OSPFv3) 1181
clear ospfv3 process 1183
clear ospfv3 redistribution 1185
clear ospfv3 routes 1187
clear ospfv3 statistics 1189
cost (OSPFv3) 1191
database-filter all out (OSPFv3) 1193
dead-interval (OSPFv3) 1195
default-cost (OSPFv3) 1197
default-information originate (OSPFv3) 1199
default-metric (OSPFv3) 1201
demand-circuit (OSPFv3) 1203
distance ospfv3 1205
distribute-list prefix-list in 1207
distribute-list prefix-list out 1209
domain-id (OSPFv3) 1212
encryption 1214
flood-reduction (OSPFv3) 1216
graceful-restart (OSPFv3) 1218
hello-interval (OSPFv3) 1220
instance (OSPFv3) 1222
interface (OSPFv3) 1224
log adjacency changes (OSPFv3) 1226
maximum interfaces (OSPFv3) 1228
maximum paths (OSPFv3) 1230

| | |
|---|------|
| maximum redistributed-prefixes (OSPFv3) | 1232 |
| mtu-ignore (OSPFv3) | 1234 |
| neighbor (OSPFv3) | 1236 |
| network (OSPFv3) | 1238 |
| nssa (OSPFv3) | 1240 |
| nsr (OSPFv3) | 1242 |
| ospfv3 name-lookup | 1244 |
| packet-size (OSPFv3) | 1245 |
| passive (OSPFv3) | 1247 |
| priority (OSPFv3) | 1249 |
| protocol shutdown (OSPFv3) | 1251 |
| range (OSPFv3) | 1252 |
| redistribute (OSPFv3) | 1254 |
| retransmit-interval (OSPFv3) | 1259 |
| router-id (OSPFv3) | 1261 |
| router ospfv3 | 1263 |
| sham-link (OSPFv3) | 1265 |
| show ospfv3 | 1267 |
| show ospfv3 border-routers | 1272 |
| show ospfv3 database | 1274 |
| show ospfv3 flood-list | 1287 |
| show ospfv3 interface | 1290 |
| show ospfv3 message-queue | 1293 |
| show ospfv3 neighbor | 1295 |
| show ospfv3 request-list | 1302 |
| show ospfv3 retransmission-list | 1305 |
| show ospfv3 routes | 1308 |
| show ospfv3 statistics rib-thread | 1311 |
| show ospfv3 summary-prefix | 1313 |
| show ospfv3 virtual-links | 1315 |
| show protocols (OSPFv3) | 1317 |
| snmp context (OSPFv3) | 1319 |
| snmp trap (OSPFv3) | 1322 |
| snmp trap rate-limit (OSPFv3) | 1323 |
| spf prefix-priority (OSPFv3) | 1325 |

stub (OSPFv3) 1327
stub-router 1329
summary-prefix (OSPFv3) 1331
timers lsa arrival 1333
timers pacing flood 1335
timers pacing lsa-group 1337
timers pacing retransmission 1339
timers throttle lsa all (OSPFv3) 1341
timers throttle spf (OSPFv3) 1343
trace (OSPFv3) 1345
transmit-delay (OSPFv3) 1347
virtual-link (OSPFv3) 1349
vrf (OSPFv3) 1351

CHAPTER 7**RIB Commands 1353**

address-family next-hop dampening disable 1355
clear route 1356
maximum prefix (RIB) 1358
lcc 1360
rcc 1362
recursion-depth-max 1364
router rib 1365
rump always-replicate 1366
show lcc statistics 1368
show rcc 1370
show rcc statistics 1372
show rcc vrf 1374
show rib 1376
show rib afi-all 1378
show rib attributes 1380
show rib client-id 1382
show rib clients 1384
show rib extcomms 1386
show rib firsthop 1388
show rib history 1390

show rib next-hop 1392
show rib opaques 1394
show rib protocols 1397
show rib recursion-depth-max 1399
show rib statistics 1401
show rib tables 1404
show rib trace 1406
show rib vpn-attributes 1408
show rib vrf 1410
show route 1412
show route backup 1418
show route best-local 1421
show route connected 1423
show route local 1425
show route longer-prefixes 1427
show route next-hop 1430
show route quarantined 1433
show route resolving-next-hop 1436
show route static 1438
show route summary 1440

CHAPTER 8**RIP Commands 1443**

authentication keychain mode (RIP) 1445
auto-summary (RIP) 1447
broadcast-for-v2 1449
clear rip 1450
clear rip database 1452
clear rip interface 1454
clear rip out-of-memory 1456
clear rip statistics 1458
default-information originate (RIP) 1460
default-metric (RIP) 1462
distance (RIP) 1464
interface (RIP) 1467
maximum-paths (RIP) 1469

- metric-zero-accept 1470
- neighbor (RIP) 1471
- nsf (RIP) 1473
- output-delay 1474
- passive-interface (RIP) 1475
- poison-reverse 1477
- receive version 1479
- redistribute (RIP) 1481
- router rip 1485
- route-policy (RIP) 1487
- send version 1489
- show protocols (RIP) 1491
- show rip 1493
- show rip database 1495
- show rip interface 1497
- show rip statistics 1504
- site-of-origin (RIP) 1506
- split-horizon disable (RIP) 1508
- timers basic 1510
- validate-update-source disable 1512
- vrf (RIP) 1513

CHAPTER 9**Routing Policy Language Commands 1515**

- abort (RPL) 1520
- add 1522
- apply 1524
- as-path in 1526
- as-path is-local 1528
- as-path length 1530
- as-path neighbor-is 1532
- as-path originates-from 1535
- as-path passes-through 1538
- as-path-set 1541
- as-path unique-length 1543
- community is-empty 1545

- community matches-any 1546
- community matches-every 1548
- community-set 1551
- delete community 1554
- delete extcommunity rt 1556
- destination in 1558
- done 1560
- drop 1562
- edit 1564
- end-global 1567
- end-policy 1568
- end-set 1570
- extcommunity rt is-empty 1572
- extcommunity rt matches-any 1574
- extcommunity rt matches-every 1576
- extcommunity rt matches-within 1578
- extcommunity-set cost 1580
- extcommunity-set rt 1582
- extcommunity-set soo 1584
- extcommunity soo is-empty 1586
- extcommunity soo matches-any 1588
- extcommunity soo matches-every 1590
- if 1592
- local-preference 1599
- med 1601
- next-hop in 1603
- orf prefix in 1605
- origin is 1607
- pass 1609
- path-type is 1611
- policy-global 1613
- prefix-set 1615
- prepend as-path 1618
- protocol 1620
- rd in 1622

rd-set 1624
replace as-path 1626
rib-has-route 1628
route-has-label 1630
route-policy (RPL) 1631
route-type is 1633
rpl editor 1635
rpl maximum 1637
set aigp-metric 1639
set community 1641
set core-tree 1643
set dampening 1645
set eigrp-metric 1647
set extcommunity cost 1649
set extcommunity rt 1651
set ip-precedence 1653
set isis-metric 1655
set label 1656
set label-mode 1657
set level 1659
set local-preference 1661
set med 1663
set metric-type (IS-IS) 1665
set metric-type (OSPF) 1667
set next-hop 1669
set origin 1671
set ospf-metric 1673
set path-selection 1674
set qos-group (RPL) 1676
set rib-metric 1678
set rip-metric 1679
set rip-tag 1681
set rpf-topology 1682
set spf-priority 1684
set tag 1686

set traffic-index 1688
set vpn-distinguisher 1690
set weight 1692
show rpl 1694
show rpl active as-path-set 1696
show rpl active community-set 1699
show rpl active extcommunity-set 1702
show rpl active prefix-set 1705
show rpl active rd-set 1708
show rpl active route-policy 1710
show rpl as-path-set 1713
show rpl as-path-set attachpoints 1715
show rpl as-path-set references 1718
show rpl community-set 1721
show rpl community-set attachpoints 1723
show rpl community-set references 1726
show rpl extcommunity-set 1729
show rpl inactive as-path-set 1732
show rpl inactive community-set 1735
show rpl inactive extcommunity-set 1738
show rpl inactive prefix-set 1741
show rpl inactive rd-set 1744
show rpl inactive route-policy 1746
show rpl maximum 1749
show rpl policy-global references 1751
show rpl prefix-set 1753
show rpl prefix-set attachpoints 1755
show rpl prefix-set references 1758
show rpl rd-set 1761
show rpl rd-set attachpoints 1763
show rpl rd-set references 1765
show rpl route-policy 1768
show rpl route-policy attachpoints 1771
show rpl route-policy inline 1774
show rpl route-policy references 1776

[show rpl route-policy uses](#) 1779
[show rpl unused as-path-set](#) 1782
[show rpl unused community-set](#) 1785
[show rpl unused extcommunity-set](#) 1788
[show rpl unused prefix-set](#) 1790
[show rpl unused rd-set](#) 1793
[show rpl unused route-policy](#) 1795
[source in](#) 1798
[suppress-route](#) 1800
[tag](#) 1802
[unsuppress-route](#) 1804
[vpn-distinguisher is](#) 1806

CHAPTER 10
Static Routing Commands 1809

[address-family \(static\)](#) 1810
[maximum path \(static\)](#) 1812
[route \(static\)](#) 1814
[router static](#) 1818
[vrf \(static\)](#) 1820

CHAPTER 11
RCMD Commands 1823

[router-convergence](#) 1824
[monitor-convergence \(IS-IS\)](#) 1826
[monitor-convergence \(OSPF\)](#) 1828
[collect-diagnostics \(RCMD\)](#) 1829
[event-buffer-size \(RCMD\)](#) 1831
[max-events-stored \(RCMD\)](#) 1833
[monitoring-interval \(RCMD\)](#) 1835
[node disable \(RCMD\)](#) 1837
[prefix-list \(monitor-convergence IS-IS\)](#) 1839
[prefix-list \(monitor-convergence OSPF\)](#) 1841
[priority \(RCMD\)](#) 1843
[protocol \(RCMD\)](#) 1845
[show rcmd isis event prefix](#) 1847
[show rcmd ospf event prefix](#) 1849

show rcmd ospf event spf 1851
storage-location 1854
track-external-routes 1856
track-summary-routes 1857

CHAPTER 12**Locator/ID Separation Protocol Commands 1859**

Locator/ID Separation Protocol on Cisco IOS XR 1860
address-family (LISP) 1861
eid-mtu 1863
eid-table 1865
etr 1867
etr accept-map-request-mapping 1869
etr map-cache-ttl 1871
etr map-server 1873
itr map-resolver 1875
locator reachability 1877
locator-set 1879
locator-table 1881
loc-reach-algorithm rloc-probing 1883
map-cache-limit 1885
map-cache 1887
map-request-source 1889
other-xtr-probe 1891
proxy-etr 1893
proxy-itr 1895
remote-rloc-probe 1897
router lisp 1899
solicit-map-request 1901
use-petr 1903



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)





Border Gateway Protocol Commands

This chapter describes the commands used to configure and monitor Border Gateway Protocol (BGP) on Cisco ASR 9000 Series Aggregation Services Routers using Cisco IOS XR software. The commands in this module configure IP Version 4 (IPv4), IP Version 6 (IPv6), Virtual Private Network Version 4 (VPNv4) routing sessions.

For detailed information about BGP concepts, configuration tasks, and examples, see the *Implementing BGP* chapter in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.



Note

Running the **show bgp** command immediately after configuring a large and complex route policy may result in timeout of the system database shown through an error message (`SYSDB-SYSDB-6-TIMEOUT_EDM`). It is recommended, that the show command be run, after the new route policy takes effect.

- [accept-own](#), page 8
- [additional-paths install backup](#), page 9
- [additional-paths receive](#), page 11
- [additional-paths selection](#), page 13
- [additional-paths send](#), page 15
- [address-family \(BGP\)](#), page 17
- [advertise best-external](#), page 21
- [advertise permanent-network](#), page 23
- [advertisement-interval](#), page 24
- [af-group](#), page 26
- [aggregate-address](#), page 28
- [aigp](#), page 31
- [aigp send-cost-community](#), page 33
- [allocate-label](#), page 35
- [allowas-in](#), page 37

- [as-format](#), page 39
- [as-override](#), page 40
- [as-path-loopcheck out disable](#), page 42
- [attribute-filter group](#), page 44
- [bfd \(BGP\)](#), page 46
- [bgp as-path-loopcheck](#), page 51
- [bgp attribute-download](#), page 52
- [bgp auto-policy-soft-reset disable](#), page 54
- [bgp bestpath as-path ignore](#), page 56
- [bgp bestpath compare-routerid](#), page 58
- [bgp bestpath cost-community ignore](#), page 60
- [bgp bestpath med always](#), page 62
- [bgp bestpath med confed](#), page 64
- [bgp bestpath med missing-as-worst](#), page 66
- [bgp client-to-client reflection disable](#), page 68
- [bgp cluster-id](#), page 70
- [bgp confederation identifier](#), page 72
- [bgp confederation peers](#), page 74
- [bgp dampening](#), page 76
- [bgp default local-preference](#), page 78
- [bgp enforce-first-as disable](#), page 79
- [bgp fast-external-fallover disable](#), page 81
- [bgp graceful-restart](#), page 82
- [bgp graceful-restart graceful-reset](#), page 84
- [bgp graceful-restart purge-time](#), page 86
- [bgp graceful-restart restart-time](#), page 88
- [bgp graceful-restart stalepath-time](#), page 90
- [bgp import-delay](#), page 92
- [bgp label-delay](#), page 94
- [bgp log neighbor changes disable](#), page 96
- [bgp maximum neighbor](#), page 98
- [bgp redistribute-internal](#), page 99
- [bgp router-id](#), page 101

- [bgp scan-time](#), page 103
- [bgp update-delay](#), page 105
- [bgp write-limit](#), page 107
- [capability additional-paths receive](#), page 109
- [capability additional-paths send](#), page 111
- [capability orf prefix](#), page 113
- [capability suppress 4-byte-as](#), page 116
- [clear bgp](#), page 119
- [clear bgp dampening](#), page 121
- [clear bgp external](#), page 123
- [clear bgp flap-statistics](#), page 125
- [clear bgp nexthop performance-statistics](#), page 127
- [clear bgp nexthop registration](#), page 129
- [clear bgp peer-drops](#), page 131
- [clear bgp performance-statistics](#), page 132
- [clear bgp self-originated](#), page 133
- [clear bgp shutdown](#), page 135
- [clear bgp soft](#), page 137
- [default-information originate \(BGP\)](#), page 140
- [default-martian-check disable](#), page 142
- [default-metric \(BGP\)](#), page 143
- [default-originate](#), page 144
- [description \(BGP\)](#), page 146
- [distance bgp](#), page 148
- [distribute bgp-ls \(ISIS\)](#), page 150
- [distribute bgp-ls \(OSPF\)](#), page 152
- [domain-distinguisher](#), page 153
- [dmz-link-bandwidth](#), page 154
- [dscp \(BGP\)](#), page 156
- [ebgp-multihop](#), page 159
- [export route-policy](#), page 161
- [export route-target](#), page 163
- [import route-policy](#), page 165

- [import route-target, page 167](#)
- [ignore-connected-check, page 169](#)
- [keychain, page 170](#)
- [keychain-disable, page 172](#)
- [keychain inheritance-disable, page 174](#)
- [label-allocation-mode, page 176](#)
- [label mode, page 178](#)
- [local-as, page 181](#)
- [maximum-paths \(BGP\), page 183](#)
- [maximum-prefix \(BGP\), page 185](#)
- [mpls activate \(BGP\), page 188](#)
- [mvpn, page 191](#)
- [neighbor \(BGP\), page 192](#)
- [neighbor-group, page 194](#)
- [network \(BGP\), page 196](#)
- [network backdoor, page 198](#)
- [next-hop-self, page 200](#)
- [next-hop-unchanged, page 203](#)
- [nexthop resolution prefix-length minimum, page 205](#)
- [nexthop route-policy, page 206](#)
- [nexthop trigger-delay, page 208](#)
- [nsr \(BGP\), page 210](#)
- [orf, page 212](#)
- [password \(BGP\), page 214](#)
- [password \(rpki-server\), page 216](#)
- [password-disable, page 218](#)
- [permanent-network, page 220](#)
- [precedence, page 222](#)
- [preference \(rpki-server\), page 224](#)
- [purge-time \(rpki-server\), page 225](#)
- [rd, page 227](#)
- [receive-buffer-size, page 229](#)
- [redistribute \(BGP\), page 231](#)

- [refresh-time \(rpki-server\), page 235](#)
- [response-time \(rpki-server\), page 237](#)
- [remote-as \(BGP\), page 239](#)
- [remove-private-as, page 242](#)
- [retain local-label, page 245](#)
- [retain route-target, page 247](#)
- [route-policy \(BGP\), page 249](#)
- [route-reflector-client, page 251](#)
- [router bgp, page 253](#)
- [rpki server, page 255](#)
- [selective-vrf-download disable, page 256](#)
- [send-buffer-size, page 258](#)
- [send-community-ebgp, page 260](#)
- [send-extended-community-ebgp, page 262](#)
- [session-group, page 264](#)
- [session-open-mode, page 266](#)
- [show bgp, page 268](#)
- [show bgp update out, page 282](#)
- [show bgp update in error process, page 284](#)
- [show bgp update out filter-group, page 285](#)
- [show bgp update out process, page 286](#)
- [show bgp update out sub-group, page 288](#)
- [show bgp update out update-group, page 290](#)
- [show bgp vrf update in error, page 292](#)
- [show bgp advertised, page 293](#)
- [show bgp af-group, page 301](#)
- [show bgp attribute-key, page 304](#)
- [show bgp cidr-only, page 309](#)
- [show bgp community, page 313](#)
- [show bgp convergence, page 318](#)
- [show bgp dampened-paths, page 321](#)
- [show bgp flap-statistics, page 326](#)
- [show bgp inconsistent-as, page 332](#)

- [show bgp labels](#), page 337
- [show bgp l2vpn rd](#), page 341
- [show bgp l2vpn vpls](#), page 344
- [show bgp neighbor-group](#), page 349
- [show bgp neighbors](#), page 353
- [show bgp neighbors nsr](#), page 375
- [show bgp nexthops](#), page 377
- [show bgp nsr](#), page 386
- [show bgp paths](#), page 390
- [show bgp policy](#), page 393
- [show bgp process](#), page 402
- [show bgp regexp](#), page 426
- [show bgp route-policy](#), page 431
- [show bgp session-group](#), page 436
- [show bgp sessions](#), page 439
- [show bgp summary](#), page 442
- [show bgp summary nsr](#), page 447
- [show bgp truncated-communities](#), page 451
- [show bgp update-group](#), page 456
- [show bgp vrf imported-routes](#), page 464
- [show protocols \(BGP\)](#), page 467
- [show svd role](#), page 470
- [show svd state](#), page 472
- [shutdown \(BGP\)](#), page 473
- [shutdown \(rpki-server\)](#), page 475
- [signalling disable](#), page 476
- [site-of-origin \(BGP\)](#), page 477
- [socket receive-buffer-size](#), page 479
- [socket send-buffer-size](#), page 481
- [soft-reconfiguration inbound](#), page 483
- [speaker-id](#), page 486
- [svd platform enable](#), page 487
- [table-policy](#), page 489

- [timers \(BGP\), page 491](#)
- [timers bgp, page 493](#)
- [transport \(rpki-server\), page 495](#)
- [ttl-security, page 497](#)
- [update limit, page 500](#)
- [update limit address-family, page 502](#)
- [update limit sub-group, page 504](#)
- [update in error-handling basic disable, page 506](#)
- [update in error-handling extended, page 508](#)
- [update out logging, page 510](#)
- [update-source, page 511](#)
- [update wait-install, page 513](#)
- [use, page 514](#)
- [username \(rpki-server\), page 519](#)
- [vrf \(BGP\), page 520](#)
- [weight, page 521](#)

accept-own

To enable handling of self-originated VPN routes containing ACCEPT_OWN community attribute, use the **accept-own** command in neighbor VPNv4 or VPNv6 address family configuration mode. To disable this functionality, either use the **no** form of this command or use the command with **inheritance-disable** keyword.

accept-own [inheritance-disable]

no accept-own

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | Disables handling of self-originated VPN routes containing ACCEPT_OWN community attribute and prevents inheritance of Accept Own from a parent configuration. |
|----------------------------|---|

Command Default

Disabled

Command Modes

Neighbor address family VPNv4

Neighbor address family VPNv6

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to enable handling of accept-own community:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#accept-own
```

additional-paths install backup



Note Effective with Release 4.0.0, the **additional-paths install backup** command was deprecated and replaced by the **additional-paths selection** command. See the [additional-paths selection, on page 13](#) command for more information.

To install a backup path into the forwarding table and provide prefix independent convergence (PIC) in case of a PE-CE link failure, use the **additional-paths install backup** command in an appropriate address family configuration mode. To prevent installing the backup path, use the **no** form of this command. To disable prefix independent convergence, use the **disable** keyword.

additional-paths install backup [disable]

no additional-paths install backup

Syntax Description

| | |
|----------------|--|
| disable | Disables installing backup path into the forwarding table. |
|----------------|--|

Command Default

None

Command Modes

VRF IPv4 address family configuration
 VRF IPv6 address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.9.0 | This command was introduced. |
| Release 4.0.0 | This command was deprecated replaced by the additional-paths selection command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable installing a backup path into the forwarding table in VPNv4 address family mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#additional-paths install backup
```

Related Commands

| Command | Description |
|--|---|
| retain local-label , on page 245 | Retains the local label until the network is converged. |

additional-paths receive

To configure receive capability of multiple paths for a prefix to the capable peers, use the **additional-paths receive** command in address-family configuration mode. To disable receive capability, use the **no** form of this command. To disable add-path receive capability for all neighbors belonging to a particular VRF address-family, use the **disable** option.

additional-paths receive [disable]

no additional-paths receive

Syntax Description

| | |
|----------------|--|
| disable | Disables advertising additional paths receive capability. Note Use the disable keyword option to disable add-path receive capability for all neighbors belonging to a specified VRF address-family. |
|----------------|--|

Syntax Description

This command has no keywords or arguments.

Command Default

None

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **additional-paths receive** command to allow add-path receive capability to be negotiated for a specified address family. When the **additional-paths receive** command is configured, the receive capability is automatically enabled for all internal BGP neighbors for a specified address family. When this command is either not configured or explicitly disabled, none of the neighbors are allowed to negotiate receive capability for the address family.

After enabling the receive capability, the session needs to be reset for the configuration to take into effect.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to enable additional paths receive capability under VPNv4 unicast address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# additional-paths receive
```

This example shows how to disable additional paths receive capability for all neighbors belonging to a particular VRF address-family (vrf1):

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)#additional-paths receive disable
```

Related Commands

| Command | Description |
|--|---|
| additional-paths send, on page 15 | Configures send capability of multiple paths for a prefix to the capable peers. |
| capability additional-paths send, on page 111 | Advertises capability of sending additional paths to the peer. |
| capability additional-paths receive, on page 109 | Advertises additional paths receive capability. |

additional-paths selection

To configure additional paths selection mode for a prefix, use the **additional-paths selection** command in address-family configuration mode. To disable the additional-paths selection mode for a prefix, use the **no** form of this command. To disable the additional-paths selection mode for a particular VRF address-family, use the **disable** option.

additional-paths selection {**route-policy** *route-policy-name*| **disable**}

no additional-paths selection route-policy *route-policy-name*

Syntax Description

| | |
|--|---|
| route-policy <i>route-policy-name</i> | Specifies the name of a route policy used for additional paths selection. |
| disable | Disables add-path selection for a particular VRF address-family. |

Command Default

None

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

To configure additional paths selection mode for some or all prefixes, use the **additional-paths selection** command by specifying a route-policy.

Use the **additional-path selection** command with an appropriate route-policy to calculate backup paths and to enable Prefix Independent Convergence (PIC) functionality. Refer *BGP Prefix Independent Convergence Unipath Primary/Backup* section in *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide* for details on the PIC functionality.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | bgp | read, write |

Examples

This example shows how to enable selection of additional paths:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router (config)# router bgp 100
RP/0/RSP0/CPU0:router (config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router (config-bgp-af)# additional-paths selection route-policy ap1
```

This example shows how to disable add-path selection for a particular VRF address-family (vrf1):

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router (config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router (config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router (config-bgp-vrf-af)#additional-paths selection disable
```

This example shows how to enable add-path selection for a particular VRF address-family (vrf2):

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router (config-bgp)#vrf vrf2
RP/0/RSP0/CPU0:router (config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router (config-bgp-vrf-af)#additional-paths selection route-policy ap2
```

additional-paths send

To configure send capability of multiple paths for a prefix to the capable peers, use the **additional-paths send** command in address-family configuration mode. To disable the send capability, use the **no** form of this command.

additional-paths send [disable]

no additional-paths send

Syntax Description

| | |
|----------------|---|
| disable | Disables advertising additional paths send capability. Note Use the disable option to disable add-path send capability for all neighbors belonging to a particular VRF address-family. |
|----------------|---|

Command Default

None

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the additional-paths send command to allow add-path send capability to be negotiated for a specified address family. When the **additional-paths send** command is configured, the send capability is automatically enabled for all internal BGP neighbors for the specified address family. When the command is either not configured or explicitly disabled, none of the neighbors are allowed to negotiate send capability for the address family.

After enabling the send capability, the session needs to be reset for the configuration to take into effect.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to enable additional paths send capability under VPNv4 4 unicast address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# additional-paths send
```

This example shows how to enable add-path selection for a particular VRF address-family (vrf1):

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config-bgp)#vrf vrf1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)#additional-paths send disable
```

Related Commands

| Command | Description |
|--|--|
| additional-paths receive, on page 11 | Configures receive capability of multiple paths for a prefix to the capable peers. |
| capability additional-paths send, on page 111 | Advertises capability of sending additional paths to the peer. |
| capability additional-paths receive, on page 109 | Advertises additional paths receive capability. |

address-family (BGP)

To enter various address family configuration modes while configuring Border Gateway Protocol (BGP), use the **address-family** command in an appropriate configuration mode. To disable support for an address family, use the **no** form of this command.

```
address-family {ipv4 {labeled-unicast| multicast| mvpn| unicast}| ipv6 {mvpn| unicast}| l2vpn vpls-vpws|
vpnv4 unicast| | link-state link-state}
```

```
no address-family
```

Syntax Description

| | |
|------------------------------|--|
| ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| ipv4 multicast | Specifies IPv4 multicast address prefixes. |
| ipv4 labeled-unicast | Specifies IPv4 labeled-unicast address prefixes. This option is available in IPv4 neighbor configuration mode and VRF neighbor configuration mode. |
| ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| vpnv4 unicast | Specifies VPN Version 4 (VPNv4) unicast address prefixes. This option is not available in VRF or VRF neighbor configuration mode. |
| l2vpn vpls-vpws | Specifies L2VPN vpls-vpws address prefixes. |
| ipv4 mvpn | Specifies IPv4 mvpn address prefixes. |
| ipv6 mvpn | specifies IPv6 mvpn address prefixes. |
| link-state link-state | Advertises link-state database of a network via BGP. |

Command Default

An address family must be explicitly configured in the router configuration mode for the address family to be active in BGP. Similarly, an address family must be configured under the neighbor for the BGP session to be established for that address family. An address family must be configured in router configuration mode before it can be configured under a neighbor.

Command Modes

Router configuration
 Neighbor configuration
 Neighbor group configuration
 VRF configuration
 VRF neighbor configuration (IPv4 address families)

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | L2VPN Address Family support was added. |
| Release 4.2.0 | The mvpn SAFI was introduced under IPv4 and IPv6. |
| Release 5.1.1 | The link-state link-state keyword was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family** command to enter various address family configuration modes while configuring BGP routing sessions. When you enter the **address-family** command from router configuration mode, you enable the address family and enter global address family configuration mode.

The IPv4 unicast address family must be configured in router configuration mode before configuring the IPv4 labeled-unicast address family for a neighbor in neighbor configuration mode. See [Table 1: Address Family Submode Support](#), on page 18.

Table 1: Address Family Submode Support

| Address Family | Supported in Router Submode | Supported in Neighbor Submode | Comments |
|----------------------|-----------------------------|-------------------------------|--|
| ipv4 unicast | yes | yes | — |
| ipv4 multicast | yes | yes | — |
| ipv4 labeled-unicast | no | yes | The ipv4 labeled-unicast address family can be configured only as a neighbor address family; however, it requires that the ipv4 unicast address family be configured as the router address family first. |
| vpnv4 unicast | yes | yes | — |
| ipv6 unicast | yes | yes | — |
| ipv6 multicast | yes | yes | — |
| l2vpn vpls-vpws | yes | yes | — |

| Address Family | Supported in Router Submode | Supported in Neighbor Submode | Comments |
|----------------|-----------------------------|-------------------------------|----------|
| ipv4 mvpn | yes | yes | — |
| ipv6 mvpn | yes | yes | — |
| link-state | yes | yes | — |

When you enter the **address-family** command from neighbor configuration mode, you activate the address family on the neighbor and enter neighbor address family configuration mode. IPv4 neighbor sessions support IPv4 unicast, multicast, labeled-unicast, and VPNv4 unicast address families. IPv6 neighbor sessions support IPv6 unicast address families.

Outbound Route Filter (ORF) capability is not supported with address-family l2vpn vpls-vpws

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to place the router in global address family configuration mode for the IPv4 address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#
```

The following example shows how to activate IPv4 multicast for neighbor 10.0.0.1 and place the router in neighbor address family configuration mode for the IPv4 multicast address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#
```

The following example shows how to place the router in global address family configuration mode for the IPv4 tunnel address family:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 12
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 tunnel
RP/0/RSP0/CPU0:router(config-bgp-af)#
```

The following example shows how to place the router in global address family link-state configuration mode:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family link-state link-state
```

```
RP/0/RSP0/CPU0:router(config-bgp-af)#
```

The following example shows how to exchange link-state information with a BGP neighbor:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router bgp 100  
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.2  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family link-state link-state  
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#
```

advertise best-external

To advertise the best-external path to the iBGP and route-reflector peers, when a locally selected bestpath is from an internal peer, use the **advertise best-external** command in an appropriate address family configuration mode. To prevent advertising the best-external path, use the **no** form of this command. To disable advertising the best-external path, use the **disable** keyword.

advertise best-external [disable]

no advertise best-external

Syntax Description

| | |
|----------------|---|
| disable | Disables best-external configuration for the VRF. |
|----------------|---|

Command Default

None

Command Modes

VRF IPv4 address family configuration
 VRF IPv6 address family configuration
 L2VPN address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration
 IPv4 address family configuration
 IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable advertising the best-external path VPNv4 unicast address family mode:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# advertise best-external
```

Related Commands

| Command | Description |
|--|---|
| additional-paths install backup, on page 9 | Installs a backup path into the forwarding table and provides prefix independent convergence (PIC) in case of a PE-CE link failure. |
| retain local-label, on page 245 | Retains the local label until the network is converged. |

advertise permanent-network

To identify the peers to whom the permanent paths must be advertised, use the **advertise permanent-network** command in the neighbor address family configuration mode. To stop advertising the permanent p, use the **no** form of this command. The permanent paths will always be advertised to peers having advertise permanent-network configuration, even if a different best-path is available. The permanent path is not advertised to peers that are not configured to receive permanent path.

The permanent path supports only prefixes in IPv4 unicast and IPv6 unicast address-families under the default Virtual Routing and Forwarding (VRF).

advertise permanent-network

no advertise permanent-network

Syntax Description This command has no arguments or keywords.

Command Modes Neighbor address-family configuration.

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 5.1.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples This example shows how to advertise permanent path:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4713
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# advertise permanent-network
```

advertisement-interval

To set the minimum interval between the sending of Border Gateway Protocol (BGP) routing updates, use the **advertisement-interval** command in an appropriate configuration mode. To remove the **advertisement-interval** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

advertisement-interval *seconds*

no advertisement-interval [*seconds*]

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Minimum interval between sending BGP routing updates (in seconds). Range is 0 to 600. |
|----------------|---|

Command Default

Default minimum interval:

For internal BGP (iBGP) peers is 0 seconds

For external BGP (eBGP) peers is 30 seconds

For customer edge (CE) peers is 0 seconds

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

VRF neighbor configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If this command configures a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to set the minimum time between sending BGP routing updates to 10 seconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 5
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RSP0/CPU0:router(config-bgp-nbr)# advertisement-interval 10
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |

af-group

To create an address family group for Border Gateway Protocol (BGP) neighbors and enter address family group configuration mode, use the **af-group** command in router configuration mode. To remove an address family group, use the **no** form of this command.

af-group *af-group-name* **address-family**

no af-group

Syntax Description

| | |
|-----------------------------|---|
| <i>af-group-name</i> | Address family group name. |
| address-family | Enters address family configuration mode. |
| ipv4 unicast | Specifies IP Version 4 (IPv4) unicast address prefixes. |
| ipv4 multicast | Specifies IPv4 multicast address prefixes. |
| ipv4 labeled-unicast | Specifies IPv4 labeled unicast address prefixes. |
| ipv4 tunnel | Specifies IPv4 tunnel address prefixes. |
| ipv6 unicast | Specifies IP Version 6 (IPv6) unicast address prefixes. |
| vpn4 unicast | Specifies VPN Version 4 (VPNv4) unicast address prefixes. |

Command Default

No BGP address family group is configured.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **af-group** command to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have address family configuration are able to use the address family group. Further, neighbors inherit the configuration parameters of the entire address family group.

You cannot define two address family groups with the same name in different address families.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to create address family group group1 and enter address family group configuration mode for IPv4 unicast. Group1 contains the next-hop-self feature, which is inherited by neighbors that use address family group1.

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# next-hop-self
```

Related Commands

| Command | Description |
|---|---|
| neighbor (BGP), on page 192 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| use, on page 514 | Inherits configuration from a neighbor group, session group, or address family group. |

aggregate-address

To create an aggregate entry in a Border Gateway Protocol (BGP) routing table, use the **aggregate-address** command in an appropriate configuration mode. To remove the **aggregate-address** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

aggregate-address *address/mask-length* [**as-set**] [**as-confed-set**] [**summary-only**] [**route-policy** *route-policy-name*]

no aggregate-address

Syntax Description

| | |
|--|--|
| <i>address</i> | Aggregate address. |
| <i>/mask-length</i> | Aggregate address mask length. |
| as-set | (Optional) Generates autonomous system set path information and community information from contributing paths. |
| as-confed-set | (Optional) Generates autonomous system confederation set path information from contributing paths. |
| summary-only | (Optional) Filters all more-specific routes from updates. |
| route-policy <i>route-policy-name</i> | (Optional) Specifies the name of a route policy used to set the attributes of the aggregate route. |

Command Default

When you do not specify this command, no aggregate entry is created in the BGP routing table.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can implement aggregate routing in BGP either by redistributing an aggregate route into BGP using the **network** command or the **aggregate-address** command.

Use the **aggregate-address** command without optional arguments to create an aggregate entry in the BGP routing table if any more-specific BGP routes are available that fall in the specified range. The aggregate route is advertised as coming from your autonomous system and has the atomic aggregate attribute set to show that information might be missing. (By default, the atomic aggregate attribute is set unless you specify the **as-set** keyword.)

Use of the **as-set** keyword creates an aggregate entry using the same rules that the command follows without this keyword. However, the advertised path for this route is an AS_SET, a set of all autonomous systems contained in all paths that are being summarized.

Do not use this form of the **aggregate-address** command when aggregating many paths because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.

Use the **as-confed-set** keyword to create an AS_CONFED_SET in the autonomous system path of the aggregate from any confederation segments in the paths being summarized. This keyword takes effect only if the **as-set** keyword is also specified.

Use of the **summary-only** keyword creates an aggregate entry (for example, 10.0.0.0/8) but suppresses advertisements of more-specific routes to all neighbors. If you want to suppress only advertisements to certain neighbors, use the **route-policy (BGP)** command in neighbor address family configuration mode with caution. If a more-specific route leaks out, all BGP speakers (the local router) prefer that route over the less-specific aggregate you generate (using longest-match routing).

Use the **route-policy** keyword to specify a routing policy for the aggregate entry. The **route-policy** keyword is used to select which more-specific information to base the aggregate entry on and which more-specific information to suppress. You can also use the keyword to modify the attributes of the aggregate entry.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to create an aggregate address. The path advertised for this route is an autonomous system set consisting of all elements contained in all paths that are being summarized.

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# aggregate-address 10.0.0.0/8 as-set
```

Related Commands

| Command | Description |
|---|---|
| network (BGP), on page 196 | Specifies the list of networks for the BGP routing process. |
| route-policy (BGP), on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor |

| Command | Description |
|--------------------|--|
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

aigp

To enable sending and receiving of accumulated interior gateway protocol (AiGP) attribute per eBGP neighbor, use the **aigp** command in global or VRF neighbor address family configuration mode. To disable this functionality, either use the **disable** keyword or use the **no** form of this command.

aigp [**disable**]

no aigp

Syntax Description

| | |
|----------------|---|
| disable | Disables sending or receiving AiGP attribute. |
|----------------|---|

Command Default

Send or receive of AiGP attribute is disabled for eBGP neighbors

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration
 Neighbor address family configuration
 VRF neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable AiGP send and receive capability under neighbor address family (IPv4 unicast):

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# aigp
```

Related Commands

| Command | Description |
|---------------------------------|-------------------------------------|
| aigp send-cost-community | Sends AiGP value in cost community. |

aigp send-cost-community

To send Accumulated Interior Gateway Protocol (AiGP) value in cost community, use the **aigp send-cost-community** command in global or VRF neighbor address family configuration mode. To disable sending AiGP value in cost community, either use the **no** form of this command or the **disable** keyword.

aigp send-cost-community {*cost-id*| **disable**} **poi** {**igp-cost**| **pre-bestpath**} [**transitive**]

no aigp send-cost-community

Syntax Description

| | |
|---------------------|--|
| <i>cost-comm-id</i> | Specifies the Cost community ID. The range is 0 to 255. |
| poi | Point of insertion for bestpath calculation. |
| igp-cost | Configures that cost community be used after iGP distance to next hop. |
| pre-bestpath | Configures cost community as first step in best path calculation. |
| transitive | (Optional) Enables transitive cost community |
| disable | Disables sending AiGP value in cost community. |

Command Default

Sending AiGP value in cost community is disabled

Command Modes

Neighbor address family configuration
VRF neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Cost community point of insertion can be configured either to be pre-bestpath or after igp cost. The **transitive** keyword is not required for iBGP sessions. However, the **transitive** keyword is required for eBGP sessions to convert AiGP metric into cost-community and advertise to the eBGP neighbors.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to enable sending AiGP value in cost community ID 254 under neighbor address family (IPv4 unicast):

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# aigp send-cost-community 254
```

Related Commands

| Command | Description |
|----------------------------------|--|
| aigp, on page 31 | Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute. |

allocate-label

To allocate Multiprotocol Label Switching (MPLS) labels for specific IPv4 unicast or VPN routing and forwarding (VRF) IPv4 unicast routes so that the BGP router can send labels with BGP routes to a neighboring router configured for labeled-unicast sessions, use the **allocate-label** command in the appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allocate-label {*route-policy route-policy-name*| **all**}

no allocate-label {*route-policy route-policy-name*| **all**}

Syntax Description

| | |
|--|--|
| all | Allocates labels for all prefixes |
| route-policy <i>route-policy-name</i> | Uses a route policy to select prefixes for label allocation. |

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration
VRF IPv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **allocate-label** command with a route policy to trigger BGP to allocate labels for all or a filtered set of global IPv4 routes (as dictated by the route policy). The command enables autonomous system border routers (ASBRs) that have labeled IPv4 unicast sessions to exchange Multiprotocol Label Switching (MPLS) labels with the IPv4 routes to the other autonomous system (AS) in Layer 3 Virtual Private Network (L3VPN) inter-AS deployments.



Note

The **allocate-label all** command is functionally equivalent to the **allocate-label route-policy route-policy-name** command when the route policy is a pass-all policy.

See *Cisco ASR 9000 Series Aggregation Services Router MPLS Configuration Guide* for information on using the **allocate-label** command for L3VPN inter-AS deployments and carrier-supporting-carrier IPv4 BGP label distribution.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable allocating labels for IPv4 routes:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# address family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# allocate-label route-policy policy_A
```

allows-in

To allow an AS path with the provider edge (PE) autonomous system number (ASN) a specified number of times, use the **allows-in** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

allows-in [*as-occurrence-number*]

no allows-in [*as-occurrence-number*]

| | | |
|---------------------------|-----------------------------|---|
| Syntax Description | <i>as-occurrence-number</i> | (Optional) Number of times a PE ASN is allowed. Range is 1 to 10. |
|---------------------------|-----------------------------|---|

Command Default No default behavior or values

Command Modes Address family group configuration
Neighbor address family configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Hub and spoke VPN networks require looping back of routing information to the hub PE through the hub customer edge (CE). See *Cisco ASR 9000 Series Aggregation Services Router MPLS Configuration Guide* for information on hub and spoke VPN networks. This looping back, in addition to the presence of the PE ASN, causes the looped-back information to be dropped by the hub PE.

The **allows-in** command prevents the looped-back information from being dropped by replacing the neighbor autonomous system number (ASN) with the PE ASN in the AS path. This allows the VPN customer to see a specified number of occurrences of the PE ASN in the AS path.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples

The following example shows how to allow five occurrences of the PE ASN:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6  
RP/0/RSP0/CPU0:router(config-bgp)# af-group group_1 address-family vpnv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# allows-in 5
```

as-format

To configure the router's Autonomous system number (ASN) notation to asdot format, use the `as-format` command in global configuration mode. To restore the system to its default condition, use the `no` form of this command.

as-format asdot

no

Syntax Description

| | |
|--------------|--|
| asdot | Specifies the Autonomous system number (ASN) notation to asdot format. |
|--------------|--|

Command Default

The default value, if the `as-format` command is not configured, is `asplain`.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the ASN notation to the asdot format:

```
RP/0/RSP0/CPU0:router(config)# as-format asdot
```

as-override

To configure a provider edge (PE) router to override the autonomous system number (ASN) of a site with the ASN of a provider, use the **as-override** command in VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

as-override [**inheritance-disable**]

no as-override [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|--|
| inheritance-disable | (Optional) Prevents the as-override command from being inherited from a parent group. |
|----------------------------|--|

Command Default

Automatic override of the ASN is disabled.

Command Modes

VRF neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-override** command in conjunction with the site-of-origin (SoO) feature, identifying the site where a route originated, and preventing routing loops between routers within a VPN.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure an ASN override:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
```

```
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr-af)# as-override
```

Related Commands

| Command | Description |
|---|--|
| site-of-origin (BGP), on page 477 | Configures the site of origin filtering. |

as-path-loopcheck out disable

To disable AS PATH loop checking for outbound updates, use the **as-path-loopcheck out disable** command in an appropriate address family configuration mode. To re-enable the default AS PATH loop checking, use the **no** form of this command.

as-path-loopcheck out disable

no as-path-loopcheck out disable

Syntax Description This command has no keywords or arguments.

Command Default AS PATH loop checking for outbound updates is enabled if there is only one neighbor and disabled if there are multiple neighbors in the update group.

Command Modes

- IPv4 address family
- IPv6 address family
- L2VPN address family
- VPNv4 address family
- VPNv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.8.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configure the **as-path-loopcheck out disable** command to disable the default behavior of PE router not announcing BGP routes to the CE router if the routes contain an AS number matching the AS number of the receiving CE router.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure **as-path-loopcheck out disable** under IPv6 unicast address family:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#router bgp 100  
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv6 unicast  
RP/0/RSP0/CPU0:router(config-bgp-af)#as-path-loopcheck out disable
```

attribute-filter group

To configure attribute-filter group command mode, use the attribute-filter group command in an appropriate configuration mode. To disable attribute-filter group command mode, use the no form of this command.

attribute-filter group *group-name*

no attribute-filter group *group-name*

Syntax Description

| | |
|-------------------|---|
| <i>group-name</i> | Specifies the name of the attribute-filter group. |
|-------------------|---|

Command Default

Attribute-filter group command mode is disabled.

Command Modes

Router configuration
Neighbor configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.3 | This command was introduced. |
| Release 5.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **attribute-filter group** command in neighbor configuration mode to configure a specific attribute filter group for a BGP neighbor.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure the attribute-filter group command mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#attribute-filter group ag_discard_med
RP/0/RSP0/CPU0:router(config-bgp-attrfg)#
```

This example shows how to configure the attribute filter group for a BGP neighbor:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.0.1.101
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 6461
RP/0/RSP0/CPU0:router(config-bgp-nbr)#update in filtering
RP/0/RSP0/CPU0:router(config-nbr-upd-filter)#attribute-filter group ag_discard_med
```

bfd (BGP)

To specify a bidirectional forwarding detection (BFD) **multiplier** and **minimum-interval** arguments per neighbor, use the **bfd** command in neighbor address family independent configuration mode. To return to the system defaults, use the **no** form of this command.

Previous to this enhancement, BFD could be configured only in global scope in BGP. This change makes available two new command-line arguments under neighbor address family independent configuration:

bfd {**multiplier** | **minimum-interval**} *value*

no bfd {**multiplier** | **minimum-interval**} *value*

Syntax Description

| | |
|--------------------------------------|--|
| multiplier <i>value</i> | Specifies the BFD session's multiplier value for the neighbor. |
| minimum-interval <i>value</i> | Specifies the BFD session's minimum-interval value for the neighbor. |

Command Default

No default per neighbor parameters are set.

Command Modes

Neighbor address family independent configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the minimum interval is changed using the **bfd minimum-interval** command, the new parameter updates all affected BFD sessions under the command mode in which the minimum interval was changed.

If the multiplier is changed using the **bfd multiplier** command, the new parameter is used to update only the BFD sessions associated with the affected neighbor gets affected.

The assumption is that when BFD fast-detect is enabled under neighbor address family independent configuration, the values for the **multiplier** and **minimum-interval** values are always derived from the per-neighbor values if they are configured; otherwise, they are to be taken from the global BGP configuration mode. In the event that this has not been explicitly stated, then these values are taken to be the default values. Also, the **bfd** arguments can be configured under neighbor-group and session-group and the inheritance adheres to the standard way of BGP configuration inheritance.

Accordingly, there are four cases in which bfd-fast detect is enabled.

This is shown in table below where the BFD value is either multiplier or minimum-interval. Local indicates per NBR value, global is the BGP global value.

| BFD value (global) | BFD value (local) | Result |
|--------------------|-------------------|---------------------|
| Yes | Yes | BFD value (local) |
| Yes | No | BFD value (global) |
| No | Yes | BFD value (local) |
| No | No | BFD value (default) |

Examples

The following example shows how to specify the BFD session's multiplier value for the neighbor:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)#neighbor 3.3.3.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 311
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd multiplier 7
RP/0/RSP0/CPU0:router(config-bgp-nbr)# neighbor 5.5.5.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd minimum-interval 318
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd multiplier 4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# vrf one
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 3.12.1.2
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# bfd minimum-interval 119
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# bfd multiplier 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# commit

RP/0/RSP0/CPU0:router# show bfd session
Interface          Dest Addr          Local det time(int*mult)  State
                   Echo              Async
-----
Gi0/2/0/2          3.3.3.2            2177ms (311ms*7)       14s (2s*7)             UP
Gi0/2/0/2.1        3.12.1.2           1190ms (119ms*10)      20s (2s*10)            UP
PO0/3/0/6          5.5.5.2            1272ms (318ms*4)       8s (2s*4)               UP

RP/0/RSP0/CPU0:router# show bfd session detail
I/f: GigabitEthernet0/2/0/2, Location: 0/2/CPU0, dest: 3.3.3.2, src: 3.3.3.1
State: UP for 0d:0h:4m:44s, number of times UP: 1
Received parameters:
Version: 1, desired tx interval: 2 s, required rx interval: 2 s
Required echo rx interval: 1 ms, multiplier: 7, diag: None
My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
Version: 1, desired tx interval: 2 s, required rx interval: 2 s
Required echo rx interval: 1 ms, multiplier: 7, diag: None
My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
Local negotiated async tx interval: 2 s
Remote negotiated async tx interval: 2 s
Desired echo tx interval: 311 ms, local negotiated echo tx interval: 311 ms
Echo detection time: 2177 ms(311 ms*7), async detection time: 14 s(2 s*7)
Local Stats:
Intervals between async packets:
Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
Last packet transmitted 313 ms ago
Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
Last packet received 1615 ms ago
Intervals between echo packets:
Tx: Number of intervals=100, min=181 ms, max=462 ms, avg=229 ms
Last packet transmitted 289 ms ago
```

```

Rx: Number of intervals=100, min=178 ms, max=461 ms, avg=229 ms
    Last packet received 287 ms ago
Latency of echo packets (time between tx and rx):
    Number of packets: 100, min=0 us, max=4 ms, avg=860 us
Session owner information:
  Client          Desired interval      Multiplier
  -----
  bgp-0          311 ms                  7

I/f: GigabitEthernet0/2/0/2.1, Location: 0/2/CPU0, dest: 3.12.1.2, src: 3.12.1.1
State: UP for 0d:0h:4m:44s, number of times UP: 1
Received parameters:
  Version: 1, desired tx interval: 2 s, required rx interval: 2 s
  Required echo rx interval: 1 ms, multiplier: 10, diag: None
  My discr: 524296, your discr: 524295, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
  Version: 1, desired tx interval: 2 s, required rx interval: 2 s
  Required echo rx interval: 1 ms, multiplier: 10, diag: None
  My discr: 524295, your discr: 524296, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
  Local negotiated async tx interval: 2 s
  Remote negotiated async tx interval: 2 s
  Desired echo tx interval: 119 ms, local negotiated echo tx interval: 119 ms
  Echo detection time: 1190 ms(119 ms*10), async detection time: 20 s(2 s*10)
Local Stats:
  Intervals between async packets:
    Tx: Number of intervals=100, min=1664 ms, max=2001 ms, avg=1838 ms
        Last packet transmitted 314 ms ago
    Rx: Number of intervals=100, min=1662 ms, max=2 s, avg=1828 ms
        Last packet received 1616 ms ago
  Intervals between echo packets:
    Tx: Number of intervals=100, min=120 ms, max=223 ms, avg=125 ms
        Last packet transmitted 112 ms ago
    Rx: Number of intervals=100, min=119 ms, max=223 ms, avg=125 ms
        Last packet received 110 ms ago
  Latency of echo packets (time between tx and rx):
    Number of packets: 100, min=0 us, max=2 ms, avg=850 us
Session owner information:
  Client          Desired interval      Multiplier
  -----
  bgp-0          119 ms                10

I/f: GigabitEthernet0/3/0/6, Location: 0/3/CPU0, dest: 5.5.5.2, src: 5.5.5.1
State: UP for 0d:0h:4m:50s, number of times UP: 1
Received parameters:
  Version: 1, desired tx interval: 2 s, required rx interval: 2 s
  Required echo rx interval: 1 ms, multiplier: 4, diag: None
  My discr: 786436, your discr: 786433, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
  Version: 1, desired tx interval: 2 s, required rx interval: 2 s
  Required echo rx interval: 1 ms, multiplier: 4, diag: None
  My discr: 786433, your discr: 786436, state UP, D/F/P/C/A: 0/0/0/1/0
Timer Values:
  Local negotiated async tx interval: 2 s
  Remote negotiated async tx interval: 2 s
  Desired echo tx interval: 318 ms, local negotiated echo tx interval: 318 ms
  Echo detection time: 1272 ms(318 ms*4), async detection time: 8 s(2 s*4)
Local Stats:
  Intervals between async packets:
    Tx: Number of intervals=100, min=1663 ms, max=2 s, avg=1821 ms
        Last packet transmitted 1740 ms ago
    Rx: Number of intervals=100, min=1663 ms, max=2001 ms, avg=1832 ms
        Last packet received 160 ms ago
  Intervals between echo packets:
    Tx: Number of intervals=100, min=181 ms, max=484 ms, avg=232 ms
        Last packet transmitted 44 ms ago
    Rx: Number of intervals=100, min=179 ms, max=484 ms, avg=232 ms
        Last packet received 41 ms ago
  Latency of echo packets (time between tx and rx):
    Number of packets: 100, min=0 us, max=3 ms, avg=540 us
Session owner information:
  Client          Desired interval      Multiplier
  -----

```

```

bgp-0          318 ms          4

RP/0/RSP0/CPU0:router# show bgp nei 3.3.3.2

BGP neighbor is 3.3.3.2
Remote AS 500, local AS 65000, external link
Remote router ID 16.0.0.1
BGP state = Established, up for 00:05:01
BFD enabled (session up): mininterval: 311 multiplier: 7
Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Neighbor capabilities:
  Route refresh: advertised and received
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 8 messages, 0 notifications, 0 in queue
Sent 9 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 30 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 2
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
  Neighbor preserved the forwarding state during latest restart
  Local restart time is 120, RIB purge time is 600 seconds
  Maximum stalepath time is 360 seconds
  Remote Restart time is 120 seconds
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is pass-all
1 accepted prefixes, 1 are bestpaths
Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
An EoR was not received during read-only mode

Connections established 1; dropped 0
Last reset 00:06:58, due to User clear requested (CEASE notification sent - administrative
reset)
Time since last notification sent to neighbor: 00:06:58
Error Code: administrative reset
Notification data sent:
  None

RP/0/RSP0/CPU0:router# show bgp nei 5.5.5.2

BGP neighbor is 5.5.5.2
Remote AS 500, local AS 65000, external link
Remote router ID 16.0.0.1
BGP state = Established, up for 00:05:04
BFD enabled (session up): mininterval: 318 multiplier: 4
Last read 00:00:58, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Neighbor capabilities:
  Route refresh: advertised and received
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 8 messages, 0 notifications, 0 in queue
Sent 9 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 30 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 2
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
  Neighbor preserved the forwarding state during latest restart
  Local restart time is 120, RIB purge time is 600 seconds
  Maximum stalepath time is 360 seconds
  Remote Restart time is 120 seconds
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is pass-all

```

```

1 accepted prefixes, 0 are bestpaths
Prefix advertised 1, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
An EoR was not received during read-only mode

Connections established 1; dropped 0
Last reset 00:07:01, due to User clear requested (CEASE notification sent - administrative
reset)
Time since last notification sent to neighbor: 00:07:01
Error Code: administrative reset
Notification data sent:
None

RP/0/RSP0/CPU0:router# show bgp vrf one nei 3.12.1.2

BGP neighbor is 3.12.1.2, vrf one
Remote AS 500, local AS 65000, external link
Remote router ID 16.0.0.1
BGP state = Established, up for 00:05:06
BFD enabled (session up): mininterval: 119 multiplier: 10
Last read 00:00:01, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Neighbor capabilities:
  Route refresh: advertised and received
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received
Received 9 messages, 0 notifications, 0 in queue
Sent 9 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 0 seconds

For Address Family: IPv4 Unicast
BGP neighbor version 2
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
  Neighbor preserved the forwarding state during latest restart
  Local restart time is 120, RIB purge time is 600 seconds
  Maximum stalepath time is 360 seconds
  Remote Restart time is 120 seconds
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is pass-all
1 accepted prefixes, 1 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
An EoR was not received during read-only mode

Connections established 1; dropped 0
Last reset 00:07:04, due to User clear requested (CEASE notification sent - administrative
reset)
Time since last notification sent to neighbor: 00:07:04
Error Code: administrative reset
Notification data sent:
None

```

bgp as-path-loopcheck

To enable loop checking in the autonomous system path of the prefixes advertised by internal Border Gateway Protocol (iBGP) peers, use the **bgp as-path-loopcheck** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp as-path-loopcheck

no bgp as-path-loopcheck

Syntax Description This command has no keywords or arguments.

Command Default When you do not specify this command, loop checking is performed only for external peers.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to configure an autonomous system path for loop checking iBGP peers:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# bgp as-path-loopcheck
```

bgp attribute-download

To enable Border Gateway Protocol (BGP) attribute download, use the **bgp attribute-download** command in an appropriate configuration mode. To disable BGP attribute download, use the **no** form of this command.

bgp attribute-download

no bgp attribute-download

Syntax Description This command has no keywords or arguments.

Command Default BGP attribute download is not enabled.

Command Modes IPv4 unicast address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When BGP attribute download is enabled using the **bgp attribute-download** command, BGP reinstalls all routes whose attributes are not currently in the RIB. Likewise, if the user disables BGP attribute download using the no form of the command, BGP reinstalls previously installed routes with a null key, and removes the attributes from the RIB.

Use the **bgp attribute-download** command to enable the Netflow BGP data export function. When attribute download is enabled, BGP downloads the attribute information for prefixes (community, extended community, and as-path) to the Routing Information Base (RIB) and Forwarding Information Base (FIB). This enables FIB to associate the prefixes with attributes and send the Netflow statistics along with the associated attributes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows the BGP routes before and after BGP attribute download is enabled and shows how to enable BGP attribute download on BGP router 50:

```
RP/0/RSP0/CPU0:router# show route bgp
```

```
B 100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:37
B 100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:37
B 100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:37
B 100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:37
B 100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:37

RP/0/RSP0/CPU0:router(config)# router bgp 50
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp attribute-download
!
!
!
RP/0/RSP0/CPU0:router# show route bgp

B 100.0.1.0/24 [200/0] via 10.0.101.1, 00:00:01
  Attribute ID 0x2
B 100.0.2.0/24 [200/0] via 10.0.101.1, 00:00:01
  Attribute ID 0x2
B 100.0.3.0/24 [200/0] via 10.0.101.1, 00:00:01
  Attribute ID 0x2
B 100.0.4.0/24 [200/0] via 10.0.101.1, 00:00:01
  Attribute ID 0x2
B 100.0.5.0/24 [200/0] via 10.0.101.1, 00:00:01
  Attribute ID 0x2
```

bgp auto-policy-soft-reset disable

To disable an automatic soft reset of Border Gateway Protocol (BGP) peers when their configured route policy is modified, use the **bgp auto-policy-soft-reset disable** command in an appropriate configuration mode. To re-enable automatic soft reset of BGP peers, use the **no** form of this command.

bgp auto-policy-soft-reset disable

no bgp auto-policy-soft-reset disable

Syntax Description This command has no keywords or arguments.

Command Default Automatic soft reset of peers is enabled.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note If the inbound policy changes, it is not always possible to perform a soft reset. This is the case if the neighbor does not support route refresh and soft-reconfiguration inbound is not configured for the neighbor. In such instances, a message is logged in the system log indicating that a manual hard reset is needed.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to disable an automatic soft reset of BGP peers when their configured route policy is modified:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
```

```
RP/0/RSP0/CPU0:router(config-bgp)# bgp auto-policy-soft-reset disable
```

bgp bestpath as-path ignore

To ignore the autonomous system path length when calculating preferred paths, use the **bgp bestpath as-path ignore** command in an appropriate configuration mode. To return the software to the default state in which it considers the autonomous system path length when calculating preferred paths, use the **no** form of this command.

bgp bestpath as-path ignore

no bgp bestpath as-path ignore

Syntax Description This command has no keywords or arguments.

Command Default The autonomous system path length is used (not ignored) when a best path is selected.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath as-path ignore** command to ignore the length of autonomous system paths when the software selects a preferred path. When the best path is selected, if this command is specified, all steps are performed as usual except comparison of the autonomous path length between candidate paths.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to configure the software to ignore the autonomous system length when performing best-path selection:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath as-path ignore
```

Related Commands

| Command | Description |
|---|--|
| bgp bestpath compare-routerid, on page 58 | Compares identical routes received from eBGP peers during the best-path selection process and selects the route with the lowest router ID. |
| bgp bestpath med always, on page 62 | Allows the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| bgp bestpath med confed, on page 64 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 66 | Enables the software to consider a missing MED attribute in a path as having a value of infinity. |

bgp bestpath compare-routerid

To compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID, use the **bgp bestpath compare-routerid** command in an appropriate configuration mode. To disable comparing identical routes received from eBGP peers during best-path selection, use the **no** form of this command.

bgp bestpath compare-routerid

no bgp bestpath compare-routerid

Syntax Description This command has no keywords or arguments.

Command Default The software does not select a new best path if it is the same as the current best path (according to the BGP selection algorithm) except for the router ID.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath compare-routerid** command to affect how the software selects the best path, in the case where there are two paths of equal cost according to the BGP selection algorithm. This command is used to force the software to select the path with the lower router ID as the best path. If this command is not used, the software continues to use whichever path is currently the best path, regardless of which has the lower router ID.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to configure the BGP speaker in autonomous system 500 to compare the router IDs of similar paths:

```
RP/0/RSP0/CPU0:router(config)# router bgp 500
```

```
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath compare-routerid
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp bestpath cost-community ignore

To configure a router that is running the Border Gateway Protocol (BGP) to not evaluate the cost community attribute during the best-path selection process, use the **bgp bestpath cost-community ignore** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

bgp bestpath cost-community ignore

no bgp bestpath cost-community ignore

Syntax Description

This command has no keywords or arguments.

Command Default

The behavior of this command is enabled by default until the cost community attribute is manually configured.

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp bestpath cost-community ignore** command to disable the evaluation of the cost community attribute to help isolate problems and troubleshoot issues that relate to BGP path selection. This command can also be used to delay the activation of cost community attribute evaluation so that cost community filtering can be deployed in a large network at the same time.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure a router to not evaluate the cost community attribute during the best-path selection process:

```
RP/0/RSP0/CPU0:router(config)# router bgp 500
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath cost-community ignore
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp bestpath med always

To allow the comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems, use the **bgp bestpath med always** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med always

no bgp bestpath med always

Syntax Description This command has no keywords or arguments.

Command Default The software does not compare MEDs for paths from neighbors in different autonomous systems.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The MED is one of the parameters that is considered by the software when selecting the best path among many alternative paths. The software chooses the path with the lowest MED.

By default, during the best-path selection process, the software makes a MED comparison only among paths from the same autonomous system. This command changes the default behavior of the software by allowing comparison of MEDs among paths regardless of the autonomous system from which the paths are received.

When the **bgp bestpath med always** command is not enabled and distributed BGP is configured, speakers calculate partial best paths only (executes the best-path steps up to the MED comparison) and send them to BGP Routing Information Base (bRIB). bRIB calculates the final best path (executes all the steps in the best-path calculation). When the **bgp bestpath med always** command is enabled and distributed BGP is configured, speakers can compare the MED across all ASs, allowing the speaker to calculate a single best path to send it to bRIB. bRIB is the ultimate process that calculates the final best path, but when the **bgp bestpath med always** command is enabled, the speakers send a single best path instead of potentially sending multiple, partial best paths

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to configure the Border Gateway Protocol (BGP) speaker in autonomous system 100 to compare MEDs among alternative paths, regardless of the autonomous system from which the paths are received:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med always
```

Related Commands

| Command | Description |
|---|---|
| bgp bestpath med confed, on page 64 | Enables MED comparison among paths learned from confederation peers. |
| bgp bestpath med missing-as-worst, on page 66 | Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp bestpath med confed

To enable Multi Exit Discriminator (MED) comparison among paths learned from confederation peers, use the **bgp bestpath med confed** command in an appropriate configuration mode. To disable the software from considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med confed

no bgp bestpath med confed

Syntax Description This command has no keywords or arguments.

Command Default The software does not compare the MED of paths containing only confederation segments, or paths containing confederation segments followed by an AS_SET, with the MED of any other paths.

Command Modes Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the MED of the following paths is not compared with the MED of any other path:

- Paths with an empty autonomous system path
- Paths beginning with an AS_SET
- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

Use the **bgp bestpath med confed** command to affect how the following types of paths are treated in the BGP best-path algorithm:

- Paths containing only confederation segments
- Paths containing confederation segments followed by an AS_SET

The MED for paths that start with an AS_SEQUENCE or that start with confederation segments followed by an AS_SEQUENCE only is compared with the MED of other paths that share the same first autonomous system number in the autonomous system sequence (the neighbor autonomous system number). This behavior is not affected by the **bgp bestpath med confed** command.

As an example, suppose that autonomous systems 65000, 65001, 65002, and 65004 are part of a confederation, but autonomous system 1 is not. Suppose that for a particular route, the following paths exist:

- Path 1: 65000 65004, med = 2, IGP metric = 20
- Path 2: 65001 65004, med = 3, IGP metric = 10
- Path 3: 65002 1, med = 1, IGP metric = 30

If the **bgp bestpath med confed** command is enabled, the software selects path 1 as the best path because it:

- Has a lower MED than path 2
- Has a lower IGP metric than path 3

The MED is not compared with path 3 because it has an external autonomous system number (that is, an AS_SEQUENCE) in the path. If the **bgp bestpath med confed** command is not enabled, then MED is not compared between any of these paths. Consequently, the software selects path 2 as the best path because it has the lowest IGP metric.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following command shows how to enable Border Gateway Protocol (BGP) software to compare MED values for paths learned from confederation peers:

```
RP/0/RSP0/CPU0:router(config)# router bgp 210
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med confed
```

Related Commands

| Command | Description |
|--|---|
| bgp bestpath med always , on page 62 | Enables MED comparison among paths from neighbors in different autonomous systems. |
| bgp bestpath med missing-as-worst , on page 66 | Specifies that the software consider a missing MED attribute in a path as having a value of infinity, making the path without a MED value the least desirable path. |
| show bgp , on page 268 | Displays entries in the BGP routing table. |

bgp bestpath med missing-as-worst

To have the software consider a missing Multi Exit Discriminator (MED) attribute in a path as having a value of infinity, making the path without a MED value the least desirable path, use the **bgp bestpath med missing-as-worst** command in an appropriate configuration mode. To disable considering the MED attribute in comparing paths, use the **no** form of this command.

bgp bestpath med missing-as-worst

no bgp bestpath med missing-as-worst

Syntax Description This command has no keywords or arguments.

Command Default The software assigns a value of 0 to the missing MED, causing the path with the missing MED attribute to be considered as the best possible MED.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read, write |

Examples The following example shows how to direct the Border Gateway Protocol (BGP) software to consider a missing MED attribute in a path as having a value of infinity, making this path the least desirable path:

```
RP/0/RSP0/CPU0:router(config)# router bgp 210
RP/0/RSP0/CPU0:router(config-bgp)# bgp bestpath med missing-as-worst
```

Related Commands

| Command | Description |
|--|--|
| bgp bestpath med always , on page 62 | Enables MED comparison among paths from neighbors in different autonomous systems. |
| bgp bestpath med confed , on page 64 | Enables MED comparison among paths learned from confederation peers. |
| show bgp , on page 268 | Displays entries in the BGP routing table. |

bgp client-to-client reflection disable

To disable reflection of routes between route-reflection clients using a Border Gateway Protocol (BGP) route reflector, use the **bgp client-to-client reflection disable** command in address family configuration mode. To re-enable client-to-client reflection, use the **no** form of this command.

bgp client-to-client reflection [*cluster-id cluster-id*] **disable**

no bgp client-to-client reflection [*cluster-id cluster-id*] **disable**

Syntax Description This command has no keywords or arguments.

Command Default Client-to-client reflection is enabled.

Command Modes Address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the clients of a route reflector are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, route reflection is not required.

Examples In this example, the three neighbors are fully meshed, so client-to-client reflection is disabled:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp client-to-client reflection disable
RP/0/RSP0/CPU0:router(config-bgp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group rrclients
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65534
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-reflector-client
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit

RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.253.21 use neighbor-group rrclients
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.253.22 use neighbor-group rrclients
```

Related Commands

| Command | Description |
|---|---|
| bgp cluster-id, on page 70 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| route-reflector-client, on page 251 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp cluster-id

To configure the cluster ID if the Border Gateway Protocol (BGP) cluster has more than one route reflector, use the **bgp cluster-id** command in an appropriate configuration mode. To remove the cluster ID, use the **no** form of this command.

bgp cluster-id *cluster-id*

no bgp cluster-id [*cluster-id*]

Syntax Description

| | |
|------------|--|
| cluster-id | Cluster ID of this router acting as a route reflector; maximum of 4 bytes. Cluster ID can be entered either as an IP address or value. Range is 1 to 4294967295. |
|------------|--|

Command Default

A cluster ID is not configured.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Together, a route reflector and its clients form a *cluster*. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the software as the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, a cluster might have more than one route reflector. If it does, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the local router as one of the route reflectors serving the cluster. Neighbor 192.168.70.24 is assigned to the default cluster with cluster-id 1.

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
```

```
RP/0/RSP0/CPU0:router(config-bgp)# bgp cluster-id 1  
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.70.24  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65534  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

Related Commands

| Command | Description |
|---|---|
| route-reflector-client, on page 251 | Configures the router as a BGP route reflector and configures the specified neighbor as its client. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp confederation identifier

To specify a Border Gateway Protocol (BGP) confederation identifier, use the **bgp confederation identifier** command in an appropriate configuration mode. To remove the confederation identifier, use the **no** form of this command.

bgp confederation identifier *as-number*

no bgp confederation identifier [*as-number*]

Syntax Description

| | |
|-----------|---|
| as-number | Autonomous system (AS) number that internally includes multiple autonomous systems. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
|-----------|---|

Command Default

No confederation identifier is configured.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

One way to reduce the internal BGP (iBGP) mesh is to divide an autonomous system into multiple autonomous systems and group them into a single confederation. Each autonomous system is fully meshed within itself, and has a few connections to another autonomous system in the same confederation. Although the peers in different autonomous systems have external BGP (eBGP) sessions, they exchange routing information as if they are iBGP peers. Specifically, the confederation maintains the next hop and local preference information, and that allows you to retain a single Interior Gateway Protocol (IGP) for all autonomous systems. To the outside world, the confederation looks like a single autonomous system.

Use the **bgp confederation identifier** command to specify the autonomous system number for the confederation. This autonomous system number is used when BGP sessions are established with external peers in autonomous systems that are not part of the confederation.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to divide the autonomous system into autonomous systems 4001, 4002, 4003, 4004, 4005, 4006, and 4007 with the confederation identifier 5. Neighbor 10.2.3.4 is a router inside the confederation. Neighbor 172.20.16.6 is outside the routing domain confederation. To the outside world, there appears to be a single autonomous system with the number 5.

```
RP/0/RSP0/CPU0:router(config)# router bgp 4001
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation identifier 5
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4002
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4003
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4004
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4005
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4006
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 4007
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4002
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbr)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 4009
```

Related Commands

| Command | Description |
|---|---|
| bgp confederation peers, on page 74 | Configures the autonomous systems that belong to the confederation. |

bgp confederation peers

To configure the autonomous systems that belong to the confederation, use the **bgp confederation peers** command in an appropriate configuration mode. To remove the autonomous system from the confederation, use the **no** form of this command.

bgp confederation peers [*as-number*]

no bgp confederation peers [*as-number*]

Syntax Description

| | |
|------------------|--|
| <i>as-number</i> | Autonomous system (AS) numbers for Border Gateway Protocol (BGP) peers that belong to the confederation. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
|------------------|--|

Command Default

No BGP peers are identified as belonging to the confederation.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The autonomous systems specified in this command are visible internally to a confederation. Each autonomous system is fully meshed within itself. The [bgp confederation identifier](#), on page 72 command specifies the confederation to which the autonomous systems belong.

To specify multiple autonomous systems, enter BGP confederation peer configuration mode then enter one *autonomous-system-number* for each command line.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows that autonomous systems 1090 and 1093 belong to a single confederation:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1090
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers 1093
```

The following example shows that autonomous systems 1095, 1096, 1097, and 1098 belong to a single confederation:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1095
RP/0/RSP0/CPU0:router(config-bgp)# bgp confederation peers
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1096
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1097
RP/0/RSP0/CPU0:router(config-bgp-confed-peers)# 1098
```

Related Commands

| Command | Description |
|--|---|
| bgp confederation identifier, on page 72 | Specifies a BGP confederation identifier. |

bgp dampening

To enable Border Gateway Protocol (BGP) route dampening or change various BGP route dampening factors, use the **bgp dampening** command in an appropriate configuration mode. To disable route dampening and reset default values, use the **no** form of this command.

bgp dampening [*half-life* [*reuse suppress max-suppress-time*]] **route-policy** *route-policy-name*

no bgp dampening [*half-life* [*reuse suppress max-suppress-time*]] **route-policy** *route-policy-name*

Syntax Description

| | |
|---|--|
| <i>half-life</i> | (Optional) Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). Penalty reduction happens every 5 seconds. Range of the half-life period is from 1 to 45 minutes. |
| <i>reuse</i> | (Optional) Value for route reuse if the flapping route penalty decreases and falls below the reuse value. When this happens, the route is unsuppressed. The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000. |
| <i>suppress</i> | (Optional) Maximum penalty value. Suppress a route when its penalty exceeds the value specified. When this happens, the route is suppressed. Range is 1 to 20000. |
| <i>max-suppress-time</i> | (Optional) Maximum time (in minutes) a route can be suppressed. Range is 1 to 255. If the <i>half-life</i> value is allowed to default, the maximum suppress time defaults to 60 minutes. |
| route-policy <i>route-policy-name</i> | (Optional) Specifies the route policy to use to set dampening parameters. |

Command Default

Route dampening is disabled.

half-life : 15 minutes

reuse : 750

suppress : 2000

max-suppress-time : four times *half-life* value

Command Modes

IPv4 address family configuration

IPv6 address family configuration

VPNv4 address family configuration

VRF IPv4 address family configuration

VPNv6 address family configuration

VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp dampening** command without arguments to enable BGP route dampening with the default parameters. The parameters can be changed by setting them on the command line or specifying them with a routing policy.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the *half-life* value to 30 minutes, the *reuse* value to 1500, the *suppress* value to 10000, and the *max-suppress-time* to 120 minutes:

```
RP/0/RSP0/CPU0:router(config)# router bgp 50
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp dampening 30 1500 10000 120
```

Related Commands

| Command | Description |
|--|--|
| clear bgp dampening, on page 121 | Clears BGP route dampening information and unsuppresses the suppressed routes. |
| clear bgp flap-statistics, on page 125 | Clears BGP flap statistics. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |
| show bgp dampened-paths, on page 321 | Displays BGP dampened routes. |
| show bgp flap-statistics, on page 326 | Displays BGP flap statistics. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |

bgp default local-preference

To change the default local preference value, use the **bgp default local-preference** command in an appropriate configuration mode. To reset the local preference value to the default of 100, use the **no** form of this command.

bgp default local-preference *value*

no bgp default local-preference [*value*]

Syntax Description

value Local preference value. Range is 0 to 4294967295. Higher values are preferable.

Command Default

Enabled with a value of 100.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Generally, the default value of 100 allows you to easily define a particular path as less preferable than paths with no local preference attribute. The preference is sent to all networking devices in the local autonomous system.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to raise the default local preference value from the default of 100 to 200:

```
RP/0/RSP0/CPU0:router(config)# router bgp 200
RP/0/RSP0/CPU0:router(config-bgp)# bgp default local-preference 200
```

bgp enforce-first-as disable

To disable the software from enforcing the first autonomous system path (known as the AS path) of a route received from an external Border Gateway Protocol (eBGP) peer to be the same as the configured remote autonomous system, use the **bgp enforce-first-as disable** command in an appropriate configuration mode. To re-enable enforcing the first AS path of a received route from an eBGP peer to be the same as the remote autonomous system, use the **no** form of this command.

bgp enforce-first-as disable

no bgp enforce-first-as disable

Syntax Description

This command has no keywords or arguments.

Command Default

By default, the software requires the first autonomous system (in the AS path) of a route received from an eBGP peer to be the same as the remote autonomous system configured.

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the software ignores any update received from an eBGP neighbor that does not have the autonomous system configured for that neighbor at the beginning of the AS path. When configured, the command applies to all eBGP peers of the router.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows a configuration in which incoming updates from eBGP neighbors are not checked to ensure the first AS number in the AS path is the same as the configured AS number for the neighbor:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# bgp enforce-first-as disable
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp fast-external-fallover disable

To disable immediately resetting the Border Gateway Protocol (BGP) sessions of any directly adjacent external peers if the link used to reach them goes down, use the **bgp fast-external-fallover disable** command in an appropriate configuration mode. To disable this function and perform an immediate reset of BGP sessions when a link between peers is lost, use the **no** form of this command.

bgp fast-external-fallover disable

no bgp fast-external-fallover disable

Syntax Description

| | |
|----------------|--------------------------------------|
| disable | Disables BGP fast external failover. |
|----------------|--------------------------------------|

Command Default

BGP sessions of any directly adjacent external peers are immediately reset if the link used to reach them goes down.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, BGP sessions of any directly adjacent external peers are immediately reset, which allows the network to recover faster when links go down between BGP peers.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to disable the automatic resetting of BGP sessions:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# bgp fast-external-fallover disable
```

bgp graceful-restart

To enable graceful restart support, use the **bgp graceful-restart** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart

no bgp graceful-restart

Syntax Description This command has no keywords or arguments.

Command Default Graceful restart support is not enabled.

Command Modes Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp graceful-restart** command to enable graceful restart functionality on the router, and also to advertise graceful restart to neighboring routers.



Note

The **bgp graceful-restart** command with no options must be used to enable graceful restart before using the **bgp graceful-restart purge-time**, **bgp graceful-restart restart-time**, **bgp graceful-restart stalepath-time**, or **bgp graceful-restart graceful-reset** commands.

When graceful restart is enabled, the BGP graceful restart capability is negotiated with neighbors in the BGP OPEN message when the session is established. If the neighbor also advertises support for graceful restart, then graceful restart is activated for that neighbor session. If the neighbor does not advertise support for graceful restart, then graceful restart is not activated for that neighbor session even though it is enabled locally.

If you enter the **bgp graceful-restart** command after some BGP sessions are established, you must restart those sessions before graceful restart takes effect. Use the **clear bgp** command to restart sessions.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable graceful restart:

```
RP/0/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)#bgp graceful-restart
```

Related Commands

| Command | Description |
|---|---|
| bgp graceful-restart graceful-reset, on page 84 | Enables a graceful reset if configuration changes force a peer reset. |
| bgp graceful-restart purge-time, on page 86 | Defines the maximum time before stale routes are purged. |
| bgp graceful-restart restart-time, on page 88 | Defines the maximum time advertised to neighbors |
| bgp graceful-restart stalepath-time, on page 90 | Defines the maximum time to wait for the End-of-RIB message from a neighbor that has been restarted before deleting learned routes. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 402 | Displays BGP process information. |

bgp graceful-restart graceful-reset

To invoke a graceful restart when configuration changes force a peer reset, use the **bgp graceful-restart graceful-reset** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

bgp graceful-restart graceful-reset

no bgp graceful-restart graceful-reset

Syntax Description This command has no keywords or arguments.

Command Default Graceful restart is not invoked when a configuration change forces a peer reset.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before enabling graceful reset using the **bgp graceful-restart graceful-reset** command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to enable graceful reset:

```
RP/0/RSP0/CPU0:router (config) #router bgp 3
RP/0/RSP0/CPU0:router (config-bgp) # bgp graceful-restart graceful-reset
```

| Related Commands | Command | Description |
|------------------|---|-----------------------------|
| | bgp graceful-restart , on page 82 | Enables a graceful restart. |

| Command | Description |
|---|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 402 | Displays BGP process information. |

bgp graceful-restart purge-time

To specify the maximum time before stale routes are purged from the routing information base (RIB) when the local BGP process restarts, use the **bgp graceful-restart purge-time** command in an appropriate configuration mode. To set the purge timer time to its default value, use the **no** form of this command.

bgp graceful-restart purge-time *seconds*

no bgp graceful-restart purge-time *seconds*

| | | |
|---------------------------|----------------|---|
| Syntax Description | <i>seconds</i> | Maximum time before stale routes are purged. Time in seconds. Range is 0 to 6000. |
|---------------------------|----------------|---|

Command Default *seconds* : 600

Command Modes Router configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the purge time using the **bgp graceful-restart purge-time** command.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

Examples The following example shows how to change the BGP purge time to 800 seconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart purge-time 800
```

Related Commands

| Command | Description |
|--|--|
| bgp graceful-restart, on page 82 | Enables a graceful restart. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 402 | Displays BGP process information. |

bgp graceful-restart restart-time

To specify a user-predicted local BGP process maximum restart time, which is advertised to neighbors during session establishment, use the **bgp graceful-restart restart-time** command in an appropriate configuration mode. To set this restart time to its default value, use the **no** form of this command.

bgp graceful-restart restart-time *seconds*

no bgp graceful-restart restart-time *seconds*

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Maximum time advertised to neighbors. Time in seconds. Range is 1 to 4095. |
|----------------|--|

Command Default

seconds : 120

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the restart timer using the **bgp graceful-restart restart-time** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to change the BGP graceful restart time to 400 seconds:

```
RP/0/RSP0/CPU0:router(config)#router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart restart-time 400
```

Related Commands

| Command | Description |
|--|--|
| bgp graceful-restart, on page 82 | Enables a graceful restart. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 402 | Displays BGP process information. |

bgp graceful-restart stalepath-time

To specify the maximum time to wait for an End-of-RIB message after a neighbor restarts, use the **bgp graceful-restart stalepath-time** command in an appropriate configuration mode. To set the stalepath timer time to its default value, use the **no** form of this command.

bgp graceful-restart stalepath-time *seconds*

no bgp graceful-restart stalepath-time *seconds*

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Maximum wait time. Time in seconds. Range is 1 to 4095. |
|----------------|---|

Command Default

seconds : 360

Command Modes

Router configuration

Command History

| Release | Modifications |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP graceful restart must be enabled using the **bgp graceful-restart** command before setting the stalepath time using the **bgp graceful-restart stalepath-time** command.

If the stalepath time is exceeded before an End-of-RIB message is received from a neighbor, paths learned from the neighbor are purged from the BGP routing table.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to change the stalepath time to 750 seconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 3
RP/0/RSP0/CPU0:router(config-bgp)# bgp graceful-restart stalepath-time 750
```

Related Commands

| Command | Description |
|--|--|
| bgp graceful-restart, on page 82 | Enables a graceful restart. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |
| show bgp process, on page 402 | Displays BGP process information. |

bgp import-delay

To enable delay for Border Gateway Protocol (BGP) batch import processing, use the **bgp import-delay** command in an appropriate configuration mode. To disable delay in batch import processing, use the no form of this command.

bgp import-delay *seconds milliseconds*

no bgp import-delay

Syntax Description

| | |
|---------------------|---|
| <i>seconds</i> | Specifies batch import processing delay in seconds. Range is 0 to 10 seconds. |
| <i>milliseconds</i> | Specifies batch import processing delay in milliseconds. Range is 0 to 999 seconds. |

Command Default

No delay is configured.

Command Modes

Address-family VPNv4 Unicast
Address-family VPNv6 Unicast

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set delay in batch import processing as two seconds and zero milliseconds:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#bgp import-delay 2 0
```

Related Commands

| Command | Description |
|---|--|
| bgp label-delay, on page 94 | Enables delay for Border Gateway Protocol (BGP) batch label processing |

bgp label-delay

To enable delay for Border Gateway Protocol (BGP) batch label processing, use the **bgp import-delay** command in an appropriate configuration mode. To disable delay in batch import processing, use the no form of this command.

bgp label-delay *seconds milliseconds*

no bgp label-delay

Syntax Description

| | |
|---------------------|--|
| <i>seconds</i> | Specifies batch label processing delay in seconds. Range is 0 to 10 seconds. |
| <i>milliseconds</i> | Specifies batch label processing delay in milliseconds. Range is 0 to 999 seconds. |

Command Default

No delay is configured.

Command Modes

Address-family IPv4 Unicast
 Address-family IPv4 Multicast
 Address-family IPv6 Unicast
 Address-family IPv6 Multicast
 Address-family VPNv4 Unicast
 Address-family VPNv6 Unicast

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set delay in batch import processing as two seconds and zero milliseconds:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#router bgp 100  
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-af)#bgp label-delay 2 0
```

Related Commands

| Command | Description |
|--|---|
| bgp import-delay, on page 92 | Enables delay for Border Gateway Protocol (BGP) batch import processing |

bgp log neighbor changes disable

To disable logging of Border Gateway Protocol (BGP) neighbor resets, use the **bgp log neighbor changes disable** command in an appropriate configuration mode. To re-enable logging of BGP neighbor resets, use the **no** form of this command.

bgp log neighbor changes disable

no bgp log neighbor changes disable

Syntax Description This command has no keywords or arguments.

Command Default BGP neighbor changes are logged.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Logging of BGP neighbor status changes (up or down) and resets is used for troubleshooting network connectivity problems and measuring network stability. Unexpected neighbor resets might indicate high error rates or high packet loss in the network, and should be investigated.

Status change message logging does not substantially affect performance, unlike, for example, enabling per-BGP update debugging. If the UNIX syslog facility is enabled, messages are sent by the software to the UNIX host running the syslog daemon so that the messages can be stored and archived on disk. If the UNIX syslog facility is not enabled, the status change messages are kept in the internal buffer of the router, and are not stored to disk.

The neighbor status change messages are not tracked if the **bgp log neighbor changes disable** command is disabled, except for the last reset reason, which is always available as output of the **show bgp neighbors** command.

Up and down messages for BGP neighbors are logged by the software by default. Use the **bgp log neighbor changes disable** command to stop logging BGP neighbor changes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to prevent the logging of neighbor changes for BGP:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# bgp log neighbor changes disable
```

Related Commands

| Command | Description |
|---|--|
| show bgp neighbors, on page 353 | Displays information about the TCP and BGP connections to neighbors. |

bgp maximum neighbor

To control the maximum number of neighbors that can be configured on the router, use the **bgp maximum neighbor** command in an appropriate configuration mode. To set the neighbor limit to the default value, use the **no** form of this command.

bgp maximum neighbor *limit*

no maximum neighbor [*limit*]

Syntax Description

| | |
|--------------|---|
| <i>limit</i> | Maximum number of neighbors. Range is 1 to 15000. |
|--------------|---|

Command Default

Default limit is 4000

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Any attempt to configure the neighbor limit below 1 or above 15000 fails. Similarly, attempting to configure the limit below the number of neighbors currently configured fails. For example, if there are 3250 neighbors configured, you cannot set the *limit* below 3250.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | write |

Examples

The following example shows how to change the default maximum neighbor limit and set it to 1200:

```
RP/0/RSP0/CPU0:router(config)#router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# bgp maximum neighbor 1200
```

bgp redistribute-internal

To allow the redistribution of internal Border Gateway Protocol (iBGP) routes into an Interior Gateway Protocol (IGP), such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF), use the **bgp redistribute-internal** command in an appropriate configuration mode. To disable the redistribution of iBGP routes into IGPs, use the **no** form of this command.

bgp redistribute-internal

no bgp redistribute-internal

Syntax Description This command has no keywords or arguments.

Command Default By default, iBGP routes are not redistributed into IGPs.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use of the **bgp redistribute-internal** command requires the **clear route *** command to be issued to reinstall all BGP routes into the IP routing table.



Note Redistributing iBGP routes into IGPs may cause routing loops to form within an autonomous system. Use this command with caution.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to redistribute iBGP routes into OSPF:

```
RP/0/RSP0/CPU0:router(config)#router bgp 1
```

```
RP/0/RSP0/CPU0:router(config-bgp)# bgp redistribute-internal
RP/0/RSP0/CPU0:router(config-bgp)# exit
RP/0/RSP0/CPU0:router(config)# router ospf areal
RP/0/RSP0/CPU0:router(config-router)# redistribute bgp 1
RP/0/RSP0/CPU0:router(config-router)# end
RP/0/RSP0/CPU0:router# clear route *
```

Related Commands

| Command | Description |
|---|---------------------------|
| clear bgp , on page 119 * | Resets all BGP neighbors. |
| clear route * | Resets all routes. |

bgp router-id

To configure a fixed router ID for a Border Gateway Protocol (BGP)-speaking router, use the **bgp router-id** command in an appropriate configuration mode. To disable a fixed router ID, use the **no** form of this command.

bgp router-id *ip-address*

no bgp router-id [*ip-address*]

Syntax Description

| | |
|-------------------|---|
| <i>ip-address</i> | IP Version 4 (IPv4) address to use as the router ID. Normally, this should be an IPv4 address assigned to the router. |
|-------------------|---|

Command Default

If no router ID is configured in BGP, BGP attempts to use the global router ID if one is configured and available. Otherwise, BGP uses the highest IP address configured on a loopback interface.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you do not use the **bgp router-id** command to configure a router ID, an IP address is not configured on any loopback interface, and no global router ID is configured, BGP neighbors remain down.

For more details on router IDs, see the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the local router with the router ID of 192.168.70.24:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
```

```
RP/0/RSP0/CPU0:router(config-bgp)#bgp router-id 192.168.70.24
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp scan-time

To configure scanning intervals of Border Gateway Protocol (BGP)-speaking networking devices, use the **bgp scan-time** command in an appropriate configuration mode. To restore the scanning interval to its default value, use the **no** form of this command.

bgp scan-time *seconds*

no bgp scan-time *seconds*

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Scanning interval (in seconds) of BGP routing information. Range is 5 to 3600 seconds. |
|----------------|--|

Command Default

The default scanning interval is 60 seconds.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | Support was removed for all address family configuration modes. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp scan-time** command to change how frequently the software processes scanner tasks, such as conditional advertisement, dynamic MED changes, and periodic maintenance tasks.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the scanning interval to 20 seconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 64500
RP/0/RSP0/CPU0:router(config-bgp-af)# bgp scan-time 20
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |

bgp update-delay

To set the maximum initial delay for a Border Gateway Protocol (BGP)-speaking router to send the first updates, use the **bgp update-delay** command in an appropriate configuration mode. To restore the initial delay to its default value, use the **no** form of this command.

bgp update-delay *seconds* [**always**]

nobgp update-delay [*seconds*][**always**]

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Delay in seconds for the router to send the first updates. Range is 0 to 3600. |
| always | (Optional) Specifies that the router always wait for the update delay time, even if all neighbors have finished sending their initial updates sooner. |

Command Default

120 seconds

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When BGP is started, it waits a specified period of time for its neighbors to establish peering sessions and to complete sending their initial updates. After all neighbors complete their initial updates, or after the update delay timer expires, the best path is calculated for each route, and the software starts sending advertisements out to its peers. This behavior improves convergence time. If the software were to advertise a route as soon as it learned it, it would have to readvertise the route each time it learned a new path that was preferred over all previously learned paths.

Use the **bgp update-delay** command to tune the maximum time the software waits after the first neighbor is established until it starts calculating best paths and sending out advertisements.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the maximum initial delay to 240 seconds:

```
RP/0/RSP0/CPU0:router(config)#router bgp 64530  
RP/0/RSP0/CPU0:router(config-bgp)# bgp update-delay 240
```

bgp write-limit



Note The **bgp write-limit** command is deprecated in Release 4.2.0, and replaced with **update limit** commands. For more information, see the commands [update limit, on page 500](#), [update limit address-family, on page 502](#), [update limit sub-group, on page 504](#).

To modify the upper bounds on update message queue lengths or to enable desynchronization, use the **bgp write-limit** command in an appropriate configuration mode. To return the bounds to their default values and to disable desynchronization, use the **no** form of this command.

bgp write-limit *group-limit global-limit* [**desynchronize**]

no bgp write-limit [*group-limit global-limit*] [**desynchronize**]

Syntax Description

| | |
|----------------------|--|
| <i>group-limit</i> | Per-update group limit on the number of update messages the software queues. Range is 500 to 100000000. Group limit cannot be greater than the global limit. |
| <i>global-limit</i> | Global limit on the number of update messages the software queues. Range is 500 to 100000000. |
| desynchronize | (Optional) Enables desynchronization. |

Command Default

group-limit : 50,000
global-limit : 250,000
 Desynchronization is off.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.0 | This command was deprecated and replaced with the update limit command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **bgp write-limit** command to configure both a per-update group and a global limit on the number of messages the software queues when updating peers. Increasing these limits can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory use during convergence. In addition, this command can be used to enable desynchronization. Desynchronization can decrease memory use and speed up convergence for the fastest neighbors if one or more neighbors in an update group process updates significantly slower than other neighbors in the same group. However, enabling desynchronization can cause a significant degradation in overall convergence time, especially if the router is experiencing high CPU utilization. For this reason, enabling desynchronization is discouraged.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure BGP to operate with a per-update group limit of 9000 messages and a global limit of 27,000 messages:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#bgp write-limit 9000 27000
```

capability additional-paths receive

To advertise capability of receiving additional paths to the peer, use the **capability additional-paths receive** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of receiving additional paths, use the **no** form of this command.

capability additional-paths receive [disable]

no capability additional-paths receive

Syntax Description

| | |
|----------------|--|
| disable | Disables advertising capability of receiving additional paths. |
|----------------|--|

Command Default

Capability is disabled.

Command Modes

Neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **capability additional-paths receive** command to selectively enable or disable additional paths receive capability negotiation for a particular neighbor or neighbor-group or session-group. Configuring **additional-paths receive** command in global address-family mode is a pre-requisite for negotiating additional paths receive capability with the peer.

If you enter the **capability additional-paths receive** command after some BGP sessions are established, you must restart those sessions for the new configuration to take effect. Use the **clear bgp** command to restart sessions.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to advertise capability of receiving additional paths:

```
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)#capability additional-paths receive
```

Related Commands

| Command | Description |
|---|--|
| additional-paths receive, on page 11 | Configures receive capability of multiple paths for a prefix to the capable peers. |
| additional-paths send, on page 15 | Configures send capability of multiple paths for a prefix to the capable peers. |
| capability additional-paths send, on page 111 | Advertises capability of sending additional paths to the peer. |

capability additional-paths send

To advertise capability of sending additional paths to the peer, use the **capability additional-paths send** command in neighbor or neighbor-group or session-group configuration mode. To disable the capability of sending additional paths, use the **no** form of this command.

capability additional paths send [disable]

no capability additional paths send

Syntax Description

| | |
|----------------|---|
| disable | Disables advertise additional paths send capability |
|----------------|---|

Command Default

Capability is disabled.

Command Modes

Neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **capability additional paths send** command to selectively enable or disable additional paths send capability negotiation for a particular neighbor or neighbor-group or session-group. Configuring the **additional-paths send** command in global address-family mode is a pre-requisite for negotiating additional paths send capability with the peer.

You must restart the BGP sessions for the new configuration to take effect. Use the **clear bgp** command to restart sessions.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to advertise capability of sending additional paths to the peer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# capability additional-paths send
```

Related Commands

| Command | Description |
|--|--|
| additional-paths receive, on page 11 | Configures receive capability of multiple paths for a prefix to the capable peers. |
| additional-paths send, on page 15 | Configures send capability of multiple paths for a prefix to the capable peers. |
| capability additional-paths receive, on page 109 | Advertises additional paths receive capability. |

capability orf prefix

To advertise prefix list-based Outbound Route Filter (ORF) capability to the Border Gateway Protocol (BGP) peer, use the **capability orf prefix** command in an appropriate configuration mode. To remove the **capability orf prefix** command from the configuration file and restore the system to its default condition in which the software does not advertise the capability, use the **no** form of this command.

capability orf prefix {receive| send| both| none}

no capability orf prefix [receive| send| both| none]

Syntax Description

| | |
|----------------|--|
| receive | Sets the capability to receive the ORF from a specified neighbor. |
| send | Sets the capability to send the ORF to a specified neighbor. |
| both | Sets the capability to receive and send the ORF from or to a specified neighbor. |
| none | Sets the capability to no for ORF receive or send from or to a specified neighbor. |

Command Default

The routing device does not receive or send route prefix filter lists.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 IPv4 neighbor address family configuration
 VRF neighbor IPv4 address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The advertisement of the prefix list ORF capability by a BGP speaker indicates whether the speaker can send prefix lists to the specified neighbor and whether it accepts prefix lists from the neighbor. The speaker sends a prefix list if it indicated the ability to send them, and if the neighbor indicated it was willing to accept them.

Similarly, the neighbor sends a prefix list to the speaker if it indicated the ability to send them and the speaker indicated the willingness to accept them.



Note The capability orf and prefix list filter specified by orf route-policy must be explicitly configured.

If the neighbor sends a prefix list and the speaker accepts it, the speaker applies the received prefix list, plus any locally configured outbound filters, to limit its outbound routing updates to the neighbor. Increased filtering prevents unwanted routing updates between neighbors and reduces resource requirements for routing update generation and processing.

Use the **capability orf prefix** command to set whether to advertise send and receive capabilities to the specified neighbor.



Note Sending a receive capability can adversely affect performance, because updates sent to that neighbor cannot be replicated for any other neighbors.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the **capability orf prefix** command:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy orfqq
RP/0/RSP0/CPU0:router:(config-rpl)# if orf prefix in (10.0.0.0/8 ge 20) then
RP/0/RSP0/CPU0:router(config-rpl)# pass
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router:(config-rpl)# if orf prefix in (1910::16 ge 120) then
RP/0/RSP0/CPU0:router(config-rpl)# pass
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router:(config-rpl)# end-policy
RP/0/RSP0/CPU0:router(config)# router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65534
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-policy pass-all out
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# capability orf prefix both
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfqq
```

Related Commands

| Command | Description |
|---|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |

| Command | Description |
|--|--|
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| show bgp neighbors , on page 353 | Displays information about BGP neighbors. Use the received prefix-filter keywords to display information on the prefix list filter. |

capability suppress 4-byte-as

To suppress 4-byte AS capability from being advertised to the BGP peer, use the **capability suppress 4-byte-as** command in the appropriate configuration mode. To remove the **capability suppress 4-byte-as** command from the configuration and restore the system to the default condition, in which the software advertises the capability, use the **no** form of this command.

capability suppress 4-byte-as [inheritance-disable]

no capability suppress 4-byte-as

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | Prevents capability suppress 4-type-as being inherited from the parent. |
|----------------------------|---|

Command Default

4-byte-as capability is advertised to the BGP peer.

Command Modes

Neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, the software advertises the 4-byte AS capability to BGP peers. To override this default behavior, use the **capability suppress 4-byte-as** command under the command modes listed in the "Command Modes" section. If configured under the neighbor group or session group, all neighbors using the group inherit the configuration. Use the **no** option to remove the command.



Caution

The BGP session resets automatically, if the 4-byte AS capability of an existing BGP session is changed by configuring **capability suppress 4-byte-as** or **capability suppress 4-byte-as inheritance-disable**.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to configure the **capability suppress 4-byte-as** command:

```
RP/0/RSP0/CPU0:router# show bgp nei 10.3.3.3 conf
neighbor 10.3.3.3
  remote-as 65000 [n:internal]
  description PE3 []
  update-source Loopback0 [n:internal]
  address-family ipv4 unicast [n:internal]

RP/0/RSP0/CPU0:router#show bgp nei 10.3.3.3
BGP neighbor is 10.3.3.3
  Remote AS 65000, local AS 65000, internal link
  Description: PE3
  Remote router ID 10.3.3.3
  BGP state = Established, up for 1w0d
  Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 25962 messages, 0 notifications, 0 in queue
  Sent 25968 messages, 1 notifications, 0 in queue
  Minimum time between advertisement runs is 0 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 1
  Update group: 0.3
  Route refresh request: received 0, sent 0
  0 accepted prefixes, 0 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
  Threshold for warning message 75%
  An EoR was received during read-only mode

  Connections established 2; dropped 1
  Last reset 1w0d, due to BGP Notification sent: hold time expired
  Time since last notification sent to neighbor: 1w0d
  Error Code: hold time expired
  Notification data sent: None

RP/0/RSP0/CPU0:router(config)#router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.3.3.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)#capability suppress 4-byte-as
RP/0/RSP0/CPU0:router(config-bgp-nbr)#commit
RP/0/RSP0/CPU0:router(config-bgp-nbr)#end

RP/0/RSP0/CPU0:router# show bgp nei 10.3.3.3

BGP neighbor is 10.3.3.3
  Remote AS 65000, local AS 65000, internal link
  Description: PE3
  Remote router ID 10.3.3.3
  BGP state = Established, up for 00:00:16
  Last read 00:00:11, hold time is 180, keepalive interval is 60 seconds
  Precedence: internet
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv4 Unicast: advertised and received
  Capability 4-byte-as suppress is configured
  Received 25966 messages, 0 notifications, 0 in queue
```

```
Sent 25972 messages, 1 notifications, 0 in queue
Minimum time between advertisement runs is 0 seconds
```

```
For Address Family: IPv4 Unicast
BGP neighbor version 1
Update group: 0.2
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
An EoR was received during read-only mode

Connections established 3; dropped 2
Last reset 00:00:43, due to Capabilty 4-byte-as configuration changed
Time since last notification sent to neighbor: 1w0d
Error Code: hold time expired
Notification data sent: None
```

With the **inheritance-disable** keyword:

```
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.101.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# capability suppress 4-byte-as inheritance-disable

RP/0/RSP0/CPU0:router# show bgp neighbor 10.0.101.1 config
neighbor 10.0.101.1
  remote-as 1 []
  address-family ipv4 unicast []

RP/0/RSP0/CPU0:router# show bgp neighbor 10.0.101.1
BGP neighbor is 10.0.101.1
Remote AS 1, local AS 100, external link
Remote router ID 0.0.0.0
BGP state = Idle
Last read 00:00:00, hold time is 180, keepalive interval is 60 seconds
Precedence: internet
Received 0 messages, 0 notifications, 0 in queue
Sent 0 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 30 seconds
```

clear bgp

To reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp** command in EXEC mode.

```
clear bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel}| ipv6 {unicast}| all {unicast| multicast|
all| labeled-unicast| tunnel}| vpnv4 unicast| vrf {vrf-name| all} {ipv4 {unicast| labeled-unicast}| ipv6
unicast}]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast and labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address prefixes. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp** command to reset the sessions of the specified group of neighbors (hard reset); it removes the TCP connection to the neighbor, removes all routes received from the neighbor from the BGP table, and then re-establishes the session with the neighbor.

If the **graceful** keyword is specified, the routes from the neighbor are not removed from the BGP table immediately, but are marked as stale. After the session is re-established, any stale route that has not been received again from the neighbor is removed.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to hard reset neighbor 10.0.0.1:

```
RP/0/RSP0/CPU0:router# clear bgp 10.0.0.1
```

Related Commands

| Command | Description |
|--|--|
| clear bgp self-originated, on page 133 | Clears self-originated routes. |
| clear bgp soft, on page 137 | Soft resets a group of BGP neighbors. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about the TCP and BGP connections to neighbors. |

clear bgp dampening

To clear Border Gateway Protocol (BGP) route dampening information and unsuppress the suppressed routes, use the **clear bgp dampening** command in EXEC configuration mode.

clear bgp dampening

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpn4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast and labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |
| <i>ip-address</i> | (Optional) IP address of the network about which to clear dampening information. |
| <i>/mask-length</i> | (Optional) Network mask applied to the IP address. |

Command Default

If no IP address is specified, dampening information for all routes is cleared.

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Always use the **clear bgp dampening** command for an individual address-family. The **all** option for address-families with clear bgp dampening should never be used during normal functioning of the system.

For example, use

```
clear bgp ipv4 unicast dampening prefix x.x.x./y
```

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear the route dampening information for all 172.20.0.0/16 IPv4 multicast paths:

```
RP/0/RSP0/CPU0:router# clear bgp ipv4 multicast dampening 172.20.0.0/16
```

Related Commands

| Command | Description |
|---|---|
| bgp dampening , on page 76 | Enables BGP route dampening or changes various BGP route dampening factors. |
| show bgp dampened-paths , on page 321 | Displays BGP dampened routes. |

clear bgp external

To clear all Border Gateway Protocol (BGP) external peers, use the **clear bgp external** command in EXEC configuration mode.

clear bgp external

| Syntax Description | |
|---|--|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| vpn6 unicast | (Optional) Specifies VPNv6 unicast address families. |
| graceful | (Optional) Clears all external peers with a hard reset and a graceful restart. This option is available when an address family is not specified. |

Command Default No default behavior or value

Command Modes EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear all BGP external peers:

```
RP/0/RSP0/CPU0:router# clear bgp external
```

clear bgp flap-statistics

To clear Border Gateway Protocol (BGP) flap counts for a specified group of routes, use the **clear bgp flap-statistics** command in EXEC configuration mode.

clear bgp flap-statistics

Syntax Description

| | |
|--|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpn4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |
| vpn6 unicast | Specifies VPNv6 unicast address families. |
| regex <i>regex</i> | (Optional) Clears flap statistics for routes whose AS paths match the regular expression. |
| route-policy <i>route-policy-name</i> | (Optional) Clears flap statistics for the specific route policy. |
| <i>network</i> | (Optional) Network for which flap counts are to be cleared. |
| <i>/mask-length</i> | (Optional) Network mask of the network for which flap counts are to be cleared. |
| <i>ip-address</i> | (Optional) Neighbor address. Clears only flap statistics for routes received from this neighbor. |

Command Default No default behavior or value

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples The following example shows how to clear the flap count for all routes (in all address families) originating in autonomous system 1:

```
RP/0/RSP0/CPU0:router#clear bgp all all flap-statistics regexp _1$
```

The following example shows how to clear the flap count for all IPv4 unicast routes received from neighbor 172.20.1.1:

```
RP/0/RSP0/CPU0:router# clear bgp ipv4 unicast flap-statistics 172.20.1.1
```

clear bgp nexthop performance-statistics

To reset the number of received notifications and the cumulative processing time for the Border Gateway Protocol (BGP) next hop, use the **clear bgp nexthop performance-statistics** command in EXEC configuration mode.

```
clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast} vpnv6 unicast} nexthop performance-statistics
```

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | Specifies tunnel address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpnv4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp nexthop performance-statistics** command to reset the total number of notifications received from the Routing Information Base (RIB) and the cumulative next-hop processing time. The following information is cleared from the **show bgp nexthops** command output:

- Total critical notifications received
- Total noncritical notifications received
- Best path deleted after last walk
- Best path changed after last walk
- Next-hop table total number of critical and noncritical notifications (Notf) and the time of the last notification received from the RIB (LastRIB) columns (only entries that have a status of unreachable [UR])

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear next-hop performance statistics:

```
RP/0/RSP0/CPU0:router# clear bgp vrf vrf_A nexthop performance statistics
```

Related Commands

| Command | Description |
|--|--|
| show bgp nexthops, on page 377 | Displays information about the BGP next-hop notifications. |

clear bgp nexthop registration

To reregister a specified next hop with the Routing Information Base (RIB), use the **clear bgp nexthop registration** command in EXEC configuration mode.

```
clear bgp {ipv4 {unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast | all | labeled-unicast} | all {unicast | multicast | all | labeled-unicast | mdt | tunnel} | vpnv4 unicast | vrf {vrf-name | all} {ipv4 {unicast | labeled-unicast} | ipv6 unicast | vpnv6 unicast} nexthop registration nexthop-address nexthop-address
```

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled-unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | Specifies tunnel address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpnv4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |
| <i>nexthop-address</i> | Address of the next hop. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp nexthop registration** command to perform an asynchronous registration of the next hop with the RIB. The **show bgp nexthops** command output shows a critical notification as the LastRIBEvent for the next hop when the **clear bgp nexthop registration** command is used.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to reregister the next hop with the RIB:

```
RP/0/RSP0/CPU0:router# clear bgp nexthop registration 10.1.1.1
```

Related Commands

| Command | Description |
|--|--|
| show bgp nexthops, on page 377 | Displays information about the BGP next-hop notifications. |

clear bgp peer-drops

To clear the connection-dropped counter, use the **clear bgp peer-drops** command in EXEC configuration mode.

```
clear bgp peer-drops {*| ip-address}
```

| Syntax Description | |
|--------------------|--|
| * | Specifies all BGP neighbors. |
| <i>ip-address</i> | IP address of a specific network neighbor. |

Command Default No default behavior or values

Command Modes EXEC configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | execute |

Examples The following example shows how to clear the connection-dropped counter for all BGP neighbors:

```
RP/0/RSP0/CPU0:router# clear bgp peer-drops *
```

| Related Commands | Command | Description |
|------------------|---|--|
| | show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |

clear bgp performance-statistics

To clear the performance statistics for all address families, use the **clear bgp performance-statistics** command.

clear bgp [*vrf* {*vrf-name*| **all**}] **performance-statistics**

Syntax Description

| | |
|-----------------|---|
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear the performance statistics for all address families:

```
RP/0/RSP0/CPU0:router# clear bgp performance-statistics
```

clear bgp self-originated

To clear Border Gateway Protocol (BGP) routes that are self-originated, use the **clear bgp self-originated** command in EXEC configuration mode.

```
clear bgp {ipv4{unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | labeled-unicast | all} | all {unicast | multicast | labeled-unicast | all} | vpnv4 unicast | vrf {vrf-name | all} | vpnv6 unicast} self-originated
```

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpnv4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Self-originated routes are routes locally originated by the **network** command, **redistribute** command, or **aggregate-address** command.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear self-originated IPv4 routes:

```
RP/0/RSP0/CPU0:router# clear bgp ipv4 unicast self-originated
```

Related Commands

| Command | Description |
|---|---|
| aggregate-address, on page 28 | Creates an aggregate entry in a BGP routing table. |
| network (BGP), on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| redistribute (BGP), on page 231 | Redistributes routes from another routing protocol into BGP. |

clear bgp shutdown

To clear all Border Gateway Protocol (BGP) neighbors that shut down due to low memory, use the **clear bgp shutdown** command in EXEC configuration mode.

```
clear bgp {ipv4{unicast | multicast | labeled-unicast | all} | ipv6 {unicast | multicast | labeled-unicast | all} | all {unicast | multicast | labeled-unicast | all} | vpnv4 unicast | vrf {vrf-name | all} | vpnv6 unicast} shutdown
```

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpnv4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to clear all shut-down BGP neighbors:

```
RP/0/RSP0/CPU0:router# clear bgp shutdown
```

Related Commands

| Command | Description |
|---|--|
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about the TCP and BGP connections to neighbors. |

clear bgp soft

To soft reset a group of Border Gateway Protocol (BGP) neighbors, use the **clear bgp soft** command in EXEC configuration mode.

```
clear bgp {ipv4{unicast | multicast | labeled-unicast | all | tunnel | mdt} | ipv6 {unicast | multicast |
labeled-unicast | all } | all {unicast | multicast | labeled-unicast | all | tunnel | mdt} | vpnv4 unicast | vrf
{vrf-name | all} | vpnv6 unicast} {* | ip-address | as| as-number | external}soft[[in| {prefix-filter} | out]]
```

Syntax Description

| | |
|---|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| labeled-unicast | Specifies labeled unicast address prefixes. |
| all | For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | Specifies tunnel address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| all | For address family, specifies prefixes for all address families. |
| vpnv4 unicast | Specifies VPNv4 unicast address families. |
| vrf | Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |
| all | For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | For VRF, specifies IPv6 unicast address families. |
| * | Soft resets all BGP neighbors. |
| <i>ip-address</i> | IP address of the neighbor to be reset. |
| as as-number | Autonomous system (AS) number for all neighbors to be reset. Range for 2-byte numbers is 1 to 65535. Range for 4-byte numbers is 1.0 to 65535.65535. |
| external | Specifies clearing of all external peers. |

| | |
|----------------------|---|
| in | (Optional) Triggers an inbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |
| prefix-filter | (Optional) Specifies to send a new Outbound Route Filter (ORF) to the neighbor. Neighbor installs the new ORF and resends its routes. |
| out | (Optional) Triggers an outbound soft reset. If the in or out keyword is not specified, both inbound and outbound soft resets are triggered. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear bgp soft** command to trigger a soft reset of the specified address families for the specified group of neighbors. This command is useful if you change the inbound or outbound policy for the neighbors, or any other configuration that affects the sending or receiving of routing updates.

If an outbound soft reset is triggered, BGP resends all routes for the address family to the given neighbors.

If an inbound soft reset is triggered, BGP by default sends a REFRESH request to the neighbor, if the neighbor has advertised the ROUTE_REFRESH capability. To determine whether the neighbor has advertised the ROUTE_REFRESH capability, use the **show bgp neighbors** command, and look for the following line of output:

```
Received route refresh capability from peer.
```

If the neighbor does not support route refresh, but the **soft-reconfiguration inbound** command is configured for the neighbor, then BGP uses the routes cached as a result of the **soft-reconfiguration inbound** command to perform the soft reset.

If you want BGP to use the cached routes even if the neighbor supports route refresh, you can use the **always** keyword when configuring the **soft-reconfiguration inbound** command.

If the neighbor does not support route refresh and the **soft-reconfiguration inbound** command is not configured, then inbound soft reset is not possible. In this case, an error is printed.

**Note**

By default, if the configuration for an inbound or outbound route policy is changed, BGP performs an automatic soft reset. Use the **bgp auto-policy-soft-reset disable** command to disable this behavior.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | execute |

Examples

The following example shows how to trigger an inbound soft clear for IPv4 unicast routes received from neighbor 10.0.0.1:

```
RP/0/RP0/CPU0:router# clear bgp ipv4 unicast 10.0.0.1 soft in
```

Related Commands

| Command | Description |
|--|---|
| bgp auto-policy-soft-reset disable, on page 54 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp, on page 119 | Resets a group of BGP neighbors. |
| clear bgp self-originated, on page 133 | Clears self-originated routes. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp neighbors, on page 353 | Displays information about the TCP and BGP connections to neighbors. |
| soft-reconfiguration inbound, on page 483 | Configures the software to store updates received from a neighbor. |

default-information originate (BGP)

To allow origination of a default route to be redistributed into the Border Gateway Protocol (BGP) from another protocol, use the **default-information originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-information originate

no default-information originate

Syntax Description This command has no arguments or keywords.

Command Default BGP does not permit redistribution of a default route into BGP.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **redistribute** command to redistribute routes from another protocol into BGP. By default, if these routes include the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6), the default route is ignored. Use the **default-information originate** command to change this behavior so that the default route is not ignored and is redistributed into BGP along with the other routes for the protocol being redistributed.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to configure BGP to redistribute the default route into BGP:

```
RP/0/RSP0/CPU0:router(config)#router bgp 164
RP/0/RSP0/CPU0:router(config-bgp)# default-information originate
```

Related Commands

| Command | Description |
|--|--|
| redistribute (BGP) , on page 231 | Redistributes routes from another protocol into BGP. |

default-martian-check disable

To disable the Martian check on the IPv4 and IPv6 prefixes, use the **default-martian-check disable** command in the address-family configuration mode. To enable the Martian check on the IPv4 and IPv6 prefixes, use the **no** form of this command.

default-martian-check disable

no default-martian-check disable

Syntax Description This command has no keywords or arguments.

Command Default None

Command Modes IPv4 address family configuration mode.
IPv6 address family configuration mode.

| Command History | Release | Modification |
|-----------------|-------------|------------------------------|
| | Release 5.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples This example shows how to disable Martian check for an IPv4 address prefix.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-af)# default-martian-check disable
```

This example shows how to disable Martian check for an IPv6 address prefix.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv6 multicast
RP/0/RSP0/CPU0:router(config-bgp-af)# default-martian-check disable
```

default-metric (BGP)

To set default metric values for the Border Gateway Protocol (BGP), use the **default-metric** command in an appropriate configuration mode. To disable metric values, use the **no** form of this command.

default-metric *value*

no default-metric [*value*]

Syntax Description

| | |
|--------------|--|
| <i>value</i> | Default metric value appropriate for the specified routing protocol. Range is 1 to 4294967295. |
|--------------|--|

Command Default

A metric is not set.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command to set the Multi Exit Discriminator (MED) to advertise to peers for routes that do not already have a metric set (routes that were received with no MED attribute).

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the BGP default metric:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# default-metric 10
```

default-originate

To cause a Border Gateway Protocol (BGP) speaker (the local router) to send the default route 0.0.0.0/0 to a neighbor for use as a default route, use the **default-originate** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

default-originate [**inheritance-disable**| **route-policy** *route-policy-name*]

no default-originate [**inheritance-disable**| **route-policy** *route-policy-name*]

Syntax Description

inheritance-disable (Optional) Prevents the **default-originate** command characteristics from being inherited from a parent group.

route-policy *route-policy-name* (Optional) Specifies the name of a route policy. The route policy allows route 0.0.0.0 to be injected conditionally. IPv6 address family is supported.

Command Default

The default route is not advertised to BGP neighbors.

Command Modes

IPv4 neighbor address family configuration
 IPv6 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 IPv4 address family group configuration
 IPv6 address family group configuration
 VRF IPv4 neighbor address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **default-originate** command does not require the presence of the default route (0.0.0.0/0 for IPv4 or ::/0 for IPv6) in the local router. When the **default-originate** command is used with a route policy, the default route is advertised if any route in the BGP table matches the policy.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to unconditionally advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# default-originate
```

The following example shows how to advertise the route 0.0.0.0/0 to the neighbor 172.20.2.3 only if a route exists in the BGP table that matches the route policy called default-default-policy:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.2.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# default-originate route-policy
default-default-policy
```

Related Commands

| Command | Description |
|---|---|
| default-information originate (BGP) , on page 140 | Allows the default route to be redistributed into BGP from another routing protocol. |
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |

description (BGP)

To annotate a neighbor, neighbor group, VPN routing and forwarding (VRF) neighbor, or session group, use the **description** command in an appropriate configuration mode. To remove the annotation, use the **no** form of this command.

description *text*

no description [*text*]

Syntax Description

| | |
|-------------|--|
| <i>text</i> | Meaningful description or comment. Maximum of 80 characters. |
|-------------|--|

Command Default

No comment or description exists.

Command Modes

Neighbor group configuration

Neighbor configuration

Session group configuration

VRF neighbor configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **description** command to provide a description of a neighbor, neighbor group, VRF neighbor, or session group. The description is used to save user comments and does not affect software function.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the description "Our best customer" on the neighbor 192.168.13.4:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)#router bgp 65000  
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 192.168.13.4  
RP/0/RSP0/CPU0:router(config-bgp-nbr)#description Our best customer
```

distance bgp

To allow the use of external, internal, and local administrative distances that could be used to prefer one class of routes over another, use the **distance bgp** command in an appropriate configuration mode. To disable the use of administrative distances, use the **no** form of this command.

distance bgp *external-distance internal-distance local-distance*

no distance bgp [*external-distance internal-distance local-distance*]

Syntax Description

| | |
|--------------------------|--|
| <i>external-distance</i> | Administrative distance for Border Gateway Protocol (BGP) external routes. External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |
| <i>internal-distance</i> | Administrative distance for BGP internal routes. Internal routes are those routes that are learned from another BGP entity within the same autonomous system. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |
| <i>local-distance</i> | Administrative distance for BGP local routes. The <i>local-distance</i> argument applies to locally generated aggregate routes (such as the routes generated by the aggregate-address command) and backdoor routes installed in the routing table. Range is 1 to 255. Routes with a distance of 255 are not installed in the routing table. |

Command Default

external-distance : 20

internal-distance : 200

local-distance : 200

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **distance bgp** command if another protocol is known to be able to provide a better route to a node than was actually learned using external BGP, or if some internal routes should be preferred by BGP.

**Note**

Changing the administrative distance of BGP internal routes is considered risky and is not recommended. One problem that can arise is the accumulation of routing table inconsistencies, which can interfere with routing.

An administrative distance is a rating of the trustworthiness of a routing information source. Numerically, an administrative distance is an integer from 1 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows that iBGP routes are preferable to locally generated routes, so the administrative distance values are set accordingly:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#distance bgp 20 20 200
```

Related Commands

| Command | Description |
|------------------|--|
| distance (IS-IS) | Defines the administrative distance assigned to routes discovered by the IS-IS protocol. |
| distance (OSPF) | Defines OSPF route administrative distances based on route type. |

distribute bgp-ls (ISIS)

To distribute ISIS link-state data using BGP LS, use the **distribute bgp-ls** command in router configuration mode. To stop link-state distribution, use the **no** form of this command.

distribute bgp-ls [**instance-id** *value*] [**level** {**1** | **2**}] [**throttle** *time*]

no distribute bgp-ls

Syntax Description

| | |
|----------------------------------|---|
| instance-id <i>value</i> | (Optional) Specifies the instance identifier defined by the router isis command. Range is from 1 to 65535. If the instance-id is not configured, the system assigned instance-id for the ISIS process will be used. |
| level 1 2 | (Optional) Displays IS-IS link-state database for Level 1 or Level 2 independently. |
| throttle | (Optional) Specifies throttle update, in seconds. Range is from 5 to 20 seconds. |

Command Default

None

Command Modes

Router configuration.

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.1.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to distribute ISIS link-state information using BGP LS:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router isis foo  
RP/0/RSP0/CPU0:router(config-isis)# distribute bgp-ls instance-id 32 level 2 throttle 5
```

distribute bgp-ls (OSPF)

To distribute OSPFv2 and OSPFv3 link-state data using BGP LS, use the **distribute bgp-ls** command in router configuration mode. To stop link-state distribution, use the **no** form of this command.

distribute bgp-ls [*instance-id value*] [*throttle time*]

no distribute bgp-ls

Syntax Description

| | |
|---------------------------------|--|
| instance-id <i>value</i> | (Optional) Specifies the instance identifier defined by the router ospf command. Range is from 1 to 65535. If the instance-id is not configured, the system assigned instance-id for the OSPF process is used. |
| throttle | (Optional) Specifies throttle time between successive link-state advertisement (LSA) updates. Range is from 0 to 3600. |

Command Default

BGP distribution is disabled.

Command Modes

Router configuration.

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.1.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to distribute OSPF link-state information using BGP LS:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)# distribute bgp-ls instance-id 32 throttle 10
```

domain-distinguisher

To configure globally unique identifier ASN for IGP domain, use the **domain-distinguisher** command in address-family link-state configuration mode. To remove unique identifier, use the **no** form of this command.

domain-distinguisher *unique-id*

no domain-distinguisher

| | | |
|---------------------------|------------------|--|
| Syntax Description | <i>unique-id</i> | Specifies four-octet unique identifier ASN. Range is from 1 to 4294967295. |
|---------------------------|------------------|--|

| | |
|------------------------|------|
| Command Default | None |
|------------------------|------|

| | |
|----------------------|--|
| Command Modes | Address-family link-state configuration. |
|----------------------|--|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 5.1.1 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|-------------------------|---|

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | bgp | read, write |

| | |
|-----------------|--|
| Examples | This example shows how to configure a unique identifier ASN: |
|-----------------|--|

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family link-state link-state
RP/0/RSP0/CPU0:router(config-bgp-af)# domain-distinguisher 1234
```

dmz-link-bandwidth

To originate a demilitarized zone (DMZ) link bandwidth extended community for the link to an eBGP or iBGP neighbor, use the **dmz-link-bandwidth** command in an Neighbor configuration mode. To stop origination of the DMZ link bandwidth extended community, use the **no** form of this command.

dmz-link-bandwidth [**inheritance-disable**]

no dmz-link-bandwidth

| Syntax Description | |
|----------------------------|---|
| inheritance-disable | (Optional) Prevents the dmz-link-bandwidth command from being inherited from a parent group. |

Command Default BGP does not originate the DMZ link bandwidth extended community.

Command Modes Neighbor configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.1.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **dmz-link-bandwidth** command to advertise the bandwidth of links that are used to exit an autonomous system.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples This example shows how to advertise the bandwidth of links to eBGP neighbors from router bgp 1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 45.67.89.01
RP/0/RSP0/CPU0:router(config-bgp-nbr)#dmz-link-bandwidth
```

Related Commands

| Command | Description |
|---|--|
| bandwidth | Configures the bandwidth of an interface. |
| maximum-paths (BGP) , on page 183 | Controls the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group , on page 264 | Creates a session group and enters session group configuration mode. |

dscp (BGP)

To set the differentiated services code point (DSCP) value, use the **dscp** command in the appropriate configuration mode. To remove the **dscp** command from the configuration file and restore the system to its default interval values, use the no form of this command.

dscp *value*

no dscp [*value*]

Syntax Description

| | |
|--------------|--|
| <i>value</i> | Value of the DSCP. The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: default , ef , af11 , af12 , af13 , af21 , af22 , af23 , af31 , af32 , af33 , af41 , af42 , af43 , cs1 , cs2 , cs3 , cs4 , cs5 , cs6 , or cs7 . |
|--------------|--|

Command Default

No default behavior or values

Command Modes

Neighbor configuration
 Neighbor session group configuration
 Neighbor group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **dscp** command to change the minimum and maximum packet thresholds for the DSCP value.

[Table 2: dscp Default Settings, on page 156](#) lists the DSCP default settings used by the **dscp** command. The DSCP value, corresponding minimum threshold, maximum threshold, and mark probability are listed. The last row of the table (the row labeled "default") shows the default settings used for any DSCP value not specifically shown in the table.

Table 2: dscp Default Settings

| DSCP (Precedence) | Minimum Threshold | Maximum Threshold | Mark Probability |
|-------------------|-------------------|-------------------|------------------|
| af11 | 32 | 40 | 1/10 |
| af12 | 28 | 40 | 1/10 |

| DSCP (Precedence) | Minimum Threshold | Maximum Threshold | Mark Probability |
|-------------------|-------------------|-------------------|------------------|
| af13 | 24 | 40 | 1/10 |
| af21 | 32 | 40 | 1/10 |
| af22 | 28 | 40 | 1/10 |
| af23 | 24 | 40 | 1/10 |
| af31 | 32 | 40 | 1/10 |
| af32 | 28 | 40 | 1/10 |
| af33 | 24 | 40 | 1/10 |
| af41 | 32 | 40 | 1/10 |
| af42 | 28 | 40 | 1/10 |
| af43 | 24 | 40 | 1/10 |
| cs1 | 22 | 40 | 1/10 |
| cs1 | 24 | 40 | 1/10 |
| cs3 | 26 | 40 | 1/10 |
| cs4 | 28 | 40 | 1/10 |
| cs5 | 30 | 40 | 1/10 |
| cs6 | 32 | 40 | 1/10 |
| cs7 | 34 | 40 | 1/10 |
| ef | 36 | 40 | 1/10 |
| default | 20 | 40 | 1/10 |

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the DSCP value to af32:

```
RP/0/RSP0/CPU0:router(config)# router bgp 5  
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.1.1.1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 100  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# dscp af32
```

ebgp-multihop

To accept and attempt Border Gateway Protocol (BGP) connections to external peers residing on networks that are not directly connected, use the **ebgp-multihop** command in an appropriate configuration mode. To disable connections to external peers and allow only direct connections between neighbors, use the **no** form of this command.

ebgp-multihop [*ttl-value*] [**mpls**]

no ebgp-multihop [*ttl-value*] [**mpls**]

Syntax Description

| | |
|------------------|--|
| <i>ttl-value</i> | (Optional) Time-to-live (TTL) value. Range is 1 to 255 hops. |
| mpls | (Optional) Disables BGP label rewrite. |

Command Default

Default TTL value is 255.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | The mpls keyword was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ebgp-multihop** command to enable multihop peerings with external BGP neighbors. The BGP protocol states that external neighbors must be directly connected (one hop away). The software enforces this by default; however, the **ebgp-multihop** command can be used to override this behavior.

Use of the **mpls** option in the **ebgp-multihop** command prevents BGP from enabling MPLS on the peering interface and also prevents allocation of Implicit-NULL rewrite labels for nexthop addresses learned from the peer. This is useful in some scenarios in which MPLS forwarding labels to the nexthops have already been learned via BGP labeled-unicast or LDP.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to allow a BGP connection to neighbor 172.20.16.6 of up to 255 hops away:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ebgp-multihop
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |

export route-policy

To configure an export route policy, use the **export route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-policy *policy-name*

no export route-policy [*policy-name*]

Syntax Description

| | |
|--------------------|--------------------------------------|
| <i>policy-name</i> | Name of the configured route policy. |
|--------------------|--------------------------------------|

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration
Global VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **export route-policy** command to define the conditions that allow specified routes to be tagged with specified route-targets.

Task ID

| Task ID | Operations |
|-------------|-------------|
| bgp | read, write |
| ip-services | read, write |

Examples

The following example shows how to configure an export route policy:

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# export route-policy policy-A
```

Related Commands

| Command | Description |
|---|--|
| import route-policy , on page 165 | Specifies a route policy to import routes into the VRF instance. |

export route-target

To configure a VPN routing and forwarding (VRF) export route-target extended community, use the **export route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

export route-target [*as-number:nn* | *ip-address:nn*]

no export route-target [*as-number:nn* | *ip-address:nn*]

Syntax Description

| | |
|----------------------|---|
| <i>as-number:nn</i> | <p>(Optional) <i>as-number</i>—Autonomous system (AS) number of the route-target extended community.</p> <ul style="list-style-type: none"> • <i>as-number</i> <ul style="list-style-type: none"> ◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. • <i>nn</i>—32-bit number |
| <i>ip-address:nn</i> | <p>(Optional) IP address of the route-target extended community.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—32-bit IP address • <i>nn</i>—16-bit number |

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration
Global VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Export route-target extended communities are associated with prefixes when advertised to remote provider edge (PE) routers. The remote PE routers import the route-target extended communities into a VRF instance that has the import route-targets that match the exported route-target extended communities.

To specify multiple route targets, enter export route target configuration mode then enter one route target for each command line.

Task ID

| Task ID | Operations |
|-------------|-------------|
| bgp | read, write |
| ip-services | read, write |

Examples

The following example shows how to specify an export route-target:

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# export route-target 500:1
```

Related Commands

| Command | Description |
|--|------------------------------------|
| import route-target, on page 167 | Specifies the import route-target. |

import route-policy

To configure an import route policy, use the **import route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-policy *policy-name*

no import route-policy [*policy-name*]

Syntax Description

| | |
|--------------------|--------------------------------------|
| <i>policy-name</i> | Name of the configured route policy. |
|--------------------|--------------------------------------|

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration
Global VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **import route-policy** command to define the conditions that allow specified routes to be imported into the VPN routing and forwarding (VRF) instance if the routes are tagged with specified route-targets.

Task ID

| Task ID | Operations |
|-------------|-------------|
| bgp | read, write |
| ip-services | read, write |

Examples

The following example shows how to allow only policy-B to be imported to VRF:

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# import route-policy policy-B
```

Related Commands

| Command | Description |
|--|--|
| export route-policy, on page 161 | Specifies a route policy to export routes from the VRF instance. |

import route-target

To configure a VPN routing and forwarding (VRF) import route-target extended community, use the **import route-target** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

import route-target [*as-number:nn* | *ip-address:nn*]

noimport route-target [*as-number:nn* | *ip-address:nn*]

Syntax Description

| | |
|----------------------|---|
| <i>as-number:nn</i> | <p>(Optional) Autonomous system (AS) number of the route-target extended community.</p> <ul style="list-style-type: none"> • <i>as-number</i> <ul style="list-style-type: none"> ◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. • <i>nn</i> —32-bit number |
| <i>ip-address:nn</i> | <p>(Optional) IP address of the route-target extended community.</p> <ul style="list-style-type: none"> • <i>ip-address</i> —32-bit IP address • <i>nn</i> —16-bit number |

Command Default

No default behavior or values

Command Modes

Global VRF IPv4 address family configuration

Global VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **import route-target** command to specify that prefixes associated with the configured import route-target extended communities are imported into the VRF instance.

To specify multiple route targets, enter import route target configuration mode, then enter one route target for each command line.

Task ID

| Task ID | Operations |
|-------------|-------------|
| bgp | read, write |
| ip-services | read, write |

Examples

The following example shows how to specify an import route-target:

```
RP/0/RSP0/CPU0:router(config)#vrf vrf-1
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# import route-target 500:99
```

Related Commands

| Command | Description |
|--|------------------------------------|
| export route-target, on page 163 | Specifies the export route-target. |

ignore-connected-check

To enable the software to bypass the directly connected next hop check for single-hop eBGP peering, use the **ignore-connected-check** command in an appropriate configuration mode. To re-enable the directly connected next hop check, use the **no** form of this command.

ignore-connected-check [**inheritance-disable**]
no ignore-connected-check

| Syntax Description | |
|---------------------|--|
| inheritance-disable | Prevents the ignore-connected-check command from being inherited from the parent. |

Command Default Ability to bypass the directly connected next hop check is disabled.

Command Modes Neighbor configuration
 Neighbor group configuration
 Session group configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.9.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to enable ignore-connected check configuration for neighbor 10.2.3.4:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.2.3.4
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ignore-connected-check
```

keychain

To apply key chain-based authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **keychain** command in an appropriate configuration mode. To disable key chain authentication, use the **no** form of this command.

keychain *name*

no keychain [*name*]

Syntax Description

| | |
|-------------|--|
| <i>name</i> | Key chain name configured using the keychain command. The name must be a maximum of 32 alphanumeric characters. |
|-------------|--|

Command Default

When this command is not specified in the appropriate configuration mode, key chain authentication is not enabled on a TCP connection between two BGP neighbors.

Command Modes

Neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Specify a key chain to enable key chain authentication between two BGP peers. Use the **keychain** command to implement hitless key rollover for authentication.

If this command is configured for a neighbor group or a session group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor override inherited values.



Note

BGP only supports HMAC-MD5 and HMAC-SHA1-12 cryptographic algorithms.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure neighbor 172.20.1.1 to use the key chain authentication configured in the keychain_A key chain:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140  
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain keychain_A
```

Related Commands

| Command | Description |
|--|--|
| keychain-disable , on page 172 | Overrides any inherited keychain configuration from a neighbor group or session group for BGP neighbors. |

keychain-disable



Note

Effective with Release 3.9.0, the **keychain-disable** command was replaced by the **keychain inheritance-disable** command. See the [keychain inheritance-disable, on page 174](#) command for more information.

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain-disable

no keychain-disable

Syntax Description

This command has no arguments or keywords.

Command Default

Configured key chains for neighbor and session groups are inherited.

Command Modes

Neighbor configuration

Neighbor group configuration

Session group configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | This command was replaced by the keychain inheritance-disable command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)#use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain-disable
```

Related Commands

| Command | Description |
|--|---|
| keychain , on page 170 | Enables key chain authentication on a TCP connection between two BGP neighbors. |

keychain inheritance-disable

To override any inherited key chain configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **keychain inheritance-disable** command in an appropriate configuration mode. To disable overriding any inherited key chain command, use the **no** form of this command.

keychain inheritance-disable

no keychain inheritance-disable

Syntax Description This command has no arguments or keywords.

Command Default Configured key chains for neighbor and session groups are inherited.

Command Modes

- Neighbor configuration
- Neighbor group configuration
- Session group configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you specify a key chain on a neighbor group or session group, all users of the group inherit the key chain. Specifying a different **keychain** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **keychain inheritance-disable** on a neighbor that uses the group disables key chain authentication for the neighbor.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to disable key chain authentication for neighbor 172.20.1.1, preventing it from inheriting the key chain keychain_A from session group group1:

```
RP/0/RSP0/CPU0:router(config)#router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# keychain keychain_A
```

```
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# keychain inheritance-disable
```

Related Commands

| Command | Description |
|---------------------------------------|---|
| keychain, on page 170 | Enables key chain authentication on a TCP connection between two BGP neighbors. |

label-allocation-mode

To set the MPLS/VPN label allocation mode, use the **label-allocation-mode** command in VRF configuration mode. To remove the **label-allocation-mode** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

label-allocation-mode [**per-ce**]

no label-allocation-mode

Syntax Description

| | |
|---------------|--|
| per-ce | Specifies that the same label is used for all the routes advertised from a unique customer edge (CE) peer or router. |
|---------------|--|

Command Default

Per-prefix is the default label allocation mode.

Command Modes

VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.3.1 | The command was hidden. This command under global IPv6 address family configuration mode was renamed to label mode . |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Each prefix that belongs to a VRF instance is advertised with a single label, causing an additional lookup to be performed in the VRF forwarding table to determine the customer edge (CE) next hop for the packet. Use the **label-allocation-mode** command with the **per-ce** keyword to avoid the additional lookup on the PE router and conserve label space. This mode allows the PE router to allocate one label for every immediate next hop. The label is directly mapped to the next hop so there is no VRF route lookup performed during data forwarding. However, the number of labels allocated is one for each CE rather than one for each prefix.

**Note**

- The **label-allocation-mode** under the global IPv6 address family configuration mode is renamed as **label mode**, in Cisco IOS-XR Software release 4.3.1 and later releases.
- With the introduction of **label mode** command, the nexthop labels will no longer be released, when **label-allocation-mode** command with the **per-ce** keyword is unconfigured.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the label allocation mode to customer edge:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109  
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1  
RP/0/RSP0/CPU0:router(config-bgp-vrf)# label-allocation-mode per-ce
```

label mode

To set the MPLS/VPN label mode based on prefix value, use the **label mode** command in an appropriate configuration mode. To remove the **label mode** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

Use this syntax for **vrf all** configuration mode under VPN IPv4/IPv6 AF (address family) mode or global IPv6 AF configuration mode:

```
label mode {per-ce | per-vrf | route-policy}
```

```
no label mode {per-ce | per-vrf | route-policy}
```

Use this syntax for IPv4/IPv6 AF configuration mode under vrf mode:

```
label mode {per-prefix | per-ce | per-vrf | route-policy}
```

```
no label mode {per-prefix | per-ce | per-vrf | route-policy}
```

Syntax Description

| | |
|---------------------|--|
| per-ce | Specifies that the same label is used for all routes advertised from a unique customer edge (CE) peer or route. |
| per-vrf | Specifies that the same label is used for all routes advertised from a unique VRF. |
| per-prefix | Specifies that the same label is used for all routes advertised from a unique prefix. Note This keyword is applicable only for IPv4/IPv6 AF configuration mode under vrf mode. |
| route-policy | Specifies a route policy to select prefixes for setting the label mode. |

Command Default

Per-prefix label mode.



Note

If a policy attached at label-mode attachpoint evaluates to pass and a **label mode** is not explicitly set, **per-prefix** is used as the default label mode.

If a policy attached at label-mode attachpoint evaluates to a drop, **per-prefix** is used as a default label mode. If any **label mode** is set explicitly in this case, it will be ignored.

Command Modes

VPNv4 address family configuration
 VPNv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

To configure label mode at VPN-AF level and to have all the VRF AFs inherit that configuration, you must use **vrf all**, which is available under VPN-AF mode.

The inheritance rules followed are:

- **label mode** configuration under VRF-AF, overrides **label-allocation-mode** configuration under VRF and **label mode** configuration under VPN-AF.
- **label-allocation-mode** configuration under VRF, overrides **label mode** configuration under VPN-AF.
- The order of priority to determine the label mode in the configurations is:
 - 1 VRF-AF: **label mode**
 - 2 VRF: **label-allocation-mode**
 - 3 VPN-AF: **label mode**
 - 4 N/A: **per-prefix**

**Note**

Even if **label mode** is in use, **per-vrf** label is allocated for connected, aggregate, and local prefixes.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

The example shows how to configure label mode selection at VPNv4 AF level:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# vrf all
RP/0/RSP0/CPU0:router(config-bgp-af)# label mode route-policy policy_A
```

The example shows how to configure label mode selection at VRF IPv4 AF level:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 109
```

```
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-bgp-vrf)# rd 1:1
RP/0/RSP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf)# label mode route-policy policy_B
```

local-as

To allow customization of the autonomous system number for external Border Gateway Protocol (eBGP) neighbor peerings, use the **local-as** command in an appropriate configuration mode. To disable customization of local autonomous system values for eBGP neighbor peerings, use the **no** form of this command.

local-as {*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] **inheritance-disable**}

no local-as [*as-number* [**no-prepend** [**replace-as** [**dual-as**]]] **inheritance-disable**]

Syntax Description

| | |
|----------------------------|---|
| <i>as-number</i> | Valid autonomous system number. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. Cannot be the autonomous system number to which the neighbor belongs. |
| no-prepend | (Optional) Specifies that local autonomous system values are not prepended to announcements from the neighbor. |
| replace-as | (Optional) Specifies that prepend only local autonomous system values to announcements to the neighbor. |
| dual-as | (Optional) Dual-AS mode. |
| inheritance-disable | Prevents local AS from being inherited from the parent. |

Command Default

The BGP autonomous system number specified in the **router bgp** command is used, except when confederations are in use. The confederation autonomous system is used for external neighbors in an autonomous system that is not part of the confederation.

Command Modes

Neighbor configuration
VRF neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The dual-as keyword was added. The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can specify the autonomous system number the local BGP uses to peer with each neighbor. The autonomous system number specified with this command cannot be the local BGP autonomous system number (specified with the **router bgp** command) or the autonomous system number of the neighbor (specified with the **remote-as** command). This command cannot be specified for internal neighbors or for external neighbors in an autonomous system that is part of a confederation.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows BGP using autonomous system 30 for the purpose of peering with neighbor 172.20.1.1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 300
RP/0/RSP0/CPU0:router(config-bgp-nbr)# local-as 30
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |

maximum-paths (BGP)

To control the maximum number of parallel routes that Border Gateway Protocol (BGP) installs in the routing table, use the **maximum-paths** command in an appropriate configuration mode. To set the maximum number of parallel routes the software installs to the default value, use the **no** form of this command.

maximum-paths {**ebgp**|**ibgp**|**eibgp**} *maximum* [**unequal-cost**]

no maximum-paths {**ebgp**|**ibgp**|**eibgp**} [*maximum*] [**unequal-cost**]

Syntax Description

| | |
|---------------------|--|
| ebgp | Specifies external BGP multipath peers. |
| ibgp | Specifies internal BGP multipath peers. |
| eibgp | Specifies internal and external BGP multipath peers. eiBGP allows simultaneous use of internal and external paths. |
| <i>maximum</i> | Maximum number of parallel routes that BGP installs in the routing table. Range is 2 to 8 |
| unequal-cost | (Optional) Allows iBGP multipaths to have different BGP next-hop Interior Gateway Protocol (IGP) metrics. This option is available when the ibgp keyword is used. |

Command Default

One path is installed in the routing table.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-paths** command to allow the BGP protocol to install multiple paths into the routing table for each prefix. Multiple paths are installed for external peers that are from the same autonomous system and

are equal cost (according to the BGP best-path algorithm). Similarly, multiple paths are installed for internal peers that are equal cost based on the BGP best-path algorithm. The IGP metric to the BGP next hop is the same as the best-path IGP metric unless the router is configured for unequal cost iBGP multipath or eBGP multipath. See *Implementing BGP* in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide* for information on the BGP best-path algorithm.



Note The **maximum-paths** command with the **eibgp** keyword cannot be configured if the **ibgp** or **ebgp** keywords have been configured, because the **eibgp** keyword is a superset of the **ibgp** or **ebgp** keywords.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to allow a maximum of four paths to a destination installed into the IPv4 unicast routing table:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# maximum-paths ebgp 4
```

maximum-prefix (BGP)

To control how many prefixes can be received from a neighbor, use the **maximum-prefix** command in an appropriate configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum-prefix *maximum* [*threshold*] [**warning-only**] [**restart** *time-interval*]

no maximum-prefix *maximum* [*threshold*] [**warning-only**] [**restart** *time-interval*]

Syntax Description

| | |
|-------------------------------------|---|
| <i>maximum</i> | Maximum number of prefixes allowed from this neighbor. Range is 1 to 4294967295. Note When using additional-paths feature, each path with a unique path ID received from a peer is counted separately for the purpose of maximum-prefix functionality. Hence, the <i>maximum</i> value should be configured appropriately when the peer is capable of sending additional-paths. |
| <i>threshold</i> | (Optional) Integer specifying at what percentage of the <i>maximum</i> argument value the software starts to generate a warning message. Range is 1 to 100. |
| warning-only | (Optional) Instructs the software to generate a log message only when the <i>maximum</i> argument value is exceeded, and not terminate the peering. |
| restart <i>time-interval</i> | (Optional) Sets the time interval (in minutes) after which peering session should be reestablished. Configure restart time interval in minutes. Range is 1 to 65535. |

Command Default

When this command is not specified, the following defaults apply:

- IPv4 Unicast: 1048576
- IPv4 Multicast: 131072
- IPv4 Labeled-unicast: 131072
- VPNv4 Unicast: 2097152
- IPv4 MDT: 131072
- IPv4 Tunnel: 1048576
- IPv6 Unicast: 524288
- IPv6 Multicast: 131072
- IPv6 Labeled-unicast: 131072
- VPNv6 Unicast: 1048576
- L2VPN EVPN: 2097152

The default threshold when a warning message is generated is 75 percent.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 IPv6 neighbor address family configuration
 VPNv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 IPv4 tunnel neighbor address family configuration
 IPv4 tunnel neighbor group address family configuration
 IPv4 tunnel address family group configuration
 VPNv4 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.1 | The default prefix limit was increased for IPv4 unicast, IPv6 unicast, VPNv4 unicast, and VPNv6 unicast address families as: <ul style="list-style-type: none"> • IPv4 unicast: 1048576 • IPv6 unicast: 524288 • VPNv4 unicast: 2097152 |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-prefix** command to configure a maximum number of prefixes that a BGP router is allowed to receive from a neighbor. It adds another mechanism (besides routing policy) to control prefixes received from a peer.

When the number of received prefixes exceeds the maximum number configured, the software terminates the peering, by default, after sending a cease notification to the neighbor. However, if the **warning-only** keyword is configured, the software writes only a log message, but continues peering with the sender. If the peer is terminated, the peer stays down until the **clear bgp** command is issued or the **restart time-interval** option is used.

This command takes effect immediately if configured on an established neighbor unless the number of prefixes received from the neighbor already exceeds the configured limits.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

This example shows the maximum number of IP Version 6 (IPv6) unicast prefixes allowed from neighbor 192.168.40.25 set to 5000, threshold value 80%, and restart time interval 20 minutes:

```
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 192.168.40.25
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#maximum-prefix 5000 80 restart 20
```

This example shows the maximum number of IP Version 4 (IPv4) unicast prefixes allowed from the neighbor 192.168.40.24 set to 1000:

```
RP/0/RSP0/CPU0:router(config-bgp)#router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#maximum-prefix 1000
```

Related Commands

| Command | Description |
|--|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp , on page 119 | Resets a BGP connection using BGP hard or soft reconfiguration. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |

mpls activate (BGP)

To enable Multiprotocol Label Switching (MPLS) on an interface basis for ASBR and CSC configurations whenever a bgp confederation configuration is used, use the **mpls activate** command in bgp configuration mode. This is needed for InterAS (option B and C) and Carrier Supporting Carrier (CSC) configurations with confederations.

The normal InterAS and CSC configurations (without confederations) do not need to enable this.

To restore the system to its default condition, use the **no** form of this command.

mpls activate *interface id*

no mpls activate *interface id*

Syntax Description

| | |
|---------------------|------------------------|
| <i>interface id</i> | Name of the interface. |
|---------------------|------------------------|

Command Default

No default behavior or values

Command Modes

Router configuration

Neighbor configuration

IPv4 address family group configuration

VPNv4 address family group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **mpls activate** command enables MPLS on the interface specified and also adds the implicit null rewrite corresponding to the peer associated with the interface. The interface specified must be the one corresponding to the inter-AS ASBR or CSC peer.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to activate MPLS for InterAS Option B (with confederations):

```
RP/0/RSP0/CPU0:router(config)#router bgp 1
```

```
  bgp confederation peers
    2002
  !
  bgp confederation identifier 4589
  bgp router-id 3.3.3.3
  mpls activate
  interface GigabitEthernet0/1/0/0
  !
  address-family ipv4 unicast
    redistribute connected
  !
  address-family vpnv4 unicast
    retain route-target all
  !
  neighbor 10.0.0.9
    remote-as 2002
  address-family ipv4 unicast
    route-policy pass in
    route-policy pass out
  !
  address-family vpnv4 unicast
    route-policy pass in
```

The following example shows how to activate MPLS for CSC (with confederations):

```
router bgp 2002
  bgp confederation peers
    1
  !
  bgp confederation identifier 4589
  bgp router-id 4.4.4.4
  address-family ipv4 unicast
    allocate-label all
  !
  address-family vpnv4 unicast
    retain route-target all
  !
  vrf foo
    rd 1:1
    mpls activate
    interface GigabitEthernet0/1/0/2
```

```

!
address-family ipv4 unicast
 redistribute connected
 allocate-label all
!
neighbor 10.0.0.1
 remote-as 1
 address-family ipv4 unicast
!
address-family ipv4 labeled-unicast
 route-policy pass in
 route-policy pass out
!
!
!
RP/0/RSP0/CPU0:router#show mpls forwarding
Local  Outgoing  Prefix          Outgoing Next Hop      Bytes
Label  Label      or ID           Interface
Switched
-----
16000  Aggregate  foo: Per-VRF Aggr[V]  \
                                           foo                0
16001  Pop        10.0.0.0/16[V]      Gi0/1/0/2    10.0.0.1      44

RP/0/RSP0/CPU0:router#show mpls interfaces
Interface          LDP      Tunnel  Enabled
-----
GigabitEthernet0/1/0/2  No      No      Yes

```

Related Commands

| Command | Description |
|--|--|
| address-family (BGP), on page 17 | Enters address family configuration mode for configuring BGP routing sessions. |

mvpn

To enable BGP instance to connect to PIM/PIM6, use the **mvpn** command in router configuration mode. To disable BGP instance -PIM/PIM6 connection, use the **no** form of this command.

mvpn

no mvpn

Syntax Description This command has no keywords or arguments.

Command Default PIM/PIM connection is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | bgp | read, write |

Examples This example shows how to configure mvpn and enable PIM/PIM6 connection:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#mvpn
```

neighbor (BGP)

To enter neighbor configuration mode for configuring Border Gateway Protocol (BGP) routing sessions, use the **neighbor** command in an appropriate configuration mode. To delete all configuration for a neighbor and terminate peering sessions with the neighbor, use the **no** form of this command.

neighbor *ip-address*

no neighbor *ip-address*

Syntax Description

| | |
|-------------------|---|
| <i>ip-address</i> | IPv4 or IPv6 IP address of the BGP-speaking neighbor. |
|-------------------|---|

Command Default

Neighbor mode is not specified.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

From router configuration mode, you can use this command to enter neighbor configuration mode.

From neighbor configuration mode, you can enter address family configuration for the neighbor by using the **address-family** command, which allows you to configure routing sessions for IP Version 4 and IP Version 6 address prefixes.

The **neighbor** command does not cause the neighbor to be configured and does not result in a peering to be established with the neighbor. To create the neighbor, you configure a remote autonomous system number by entering the **remote-as** command, or the neighbor can inherit a remote autonomous system from a neighbor group or session group if the **use** command is applied.



Note

A neighbor must have a remote autonomous system number, and an IP address and address family must be enabled on the neighbor.

Unlike IPv4, IPv6 must be enabled before any IPv6 neighbors can be defined. Enable IPv6 in router configuration mode using the **address-family** command.

**Note**

Configuration for the neighbor cannot occur (peering is not established) until the neighbor is given a remote as-number and neighbor address.

The **no** form of this command causes the peering with the neighbor to be terminated and all configuration that relates to the neighbor to be removed.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to place the router in neighbor configuration mode for BGP routing process 1 and configure the neighbor IP address 172.168.40.24 as a BGP peer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65000
```

The following example shows how to enable IPv6 for BGP, then place the router in neighbor configuration mode for an IPv6 neighbor, 3000::1, and configure neighbor 3000::1 as a BGP peer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 3000::1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2002
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv6 unicast
```

Related Commands

| Command | Description |
|--|---|
| address-family (BGP), on page 17 | Enters address family configuration mode for configuring BGP routing sessions. |
| remote-as (BGP), on page 239 | Adds an entry to the BGP neighbor table. |
| use, on page 514 | Inherits characteristics from a neighbor group, session group, or address family group. |

neighbor-group

To create a neighbor group and enter neighbor group configuration mode, use the **neighbor-group** command in router configuration mode. To remove a neighbor group and delete all configuration associated with the group, use the **no** form of this command.

neighbor-group *name*

no neighbor-group *name*

Syntax Description

| | |
|-------------|----------------------|
| <i>name</i> | Neighbor group name. |
|-------------|----------------------|

Command Default

No neighbor group mode is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **neighbor-group** command puts the router in neighbor group configuration mode and creates a neighbor group.

A neighbor group helps you apply the same configuration to one or more neighbors. After a neighbor group is configured, each neighbor can inherit the configuration through the **use** command. If a neighbor is configured to use a neighbor group, the neighbor, by default, inherits the entire configuration of the neighbor group, which includes the address family-independent and address family-specific configurations. The inherited configuration can be overridden if you directly configure commands for the neighbor or if you configure session groups or address family groups with the **use** command.

From neighbor group configuration mode, you can configure address family-independent parameters for the neighbor group. To enter address family-specific configuration for the neighbor group, use the **address-family** command when in the neighbor group configuration mode.



Note

If an address family is configured for a neighbor group, neighbors that use the neighbor group attempt to exchange routes in that address family.

The **no** form of this command ordinarily causes all configuration for the neighbor group to be removed. If using the **no** form would result in a neighbor losing its remote autonomous system number, the configuration is rejected. In this scenario, the neighbor configuration must be either removed or configured with a remote autonomous system number before the neighbor group configuration can be removed.



Note Neighbor groups should not be configured with a mixture of IPv4 and IPv6 address families, because such a neighbor group is not usable by any neighbor. Note that within the Cisco IOS XR system configuration architecture, it is possible to create such a neighbor group; however, any attempt to use it is rejected.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to create a neighbor group called group1 that has IP Version 4 (IPv4) unicast and IPv4 multicast activated along with various configuration features. The neighbor group is used by neighbor 10.0.0.1 and neighbor 10.0.0.2, which allows them to inherit the entire group1 configuration.

```
RP/0/RSP0/CPU0:router(config)# router bgp 65530
RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# remote-as 65535
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 2
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# send-community-ebgp
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# next-hop-self
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 10.0.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|--|--|
| address-family (BGP), on page 17 | Enters various address family configuration modes for configuring BGP routing sessions. |
| neighbor (BGP), on page 192 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| use, on page 514 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

network (BGP)

To specify that the Border Gateway Protocol (BGP) routing process should originate and advertise a locally known network to its neighbors, use the **network** command in an appropriate configuration mode. To disable originating or advertising the network to neighbors, use the **no** form of this command.

network {*ip-address/prefix-length* | *ip-address mask*} [**route-policy** *route-policy-name*]

no network{*ip-address/prefix-length* | *ip-address mask*} [**route-policy** *route-policy-name*]

Syntax Description

| | |
|--|--|
| <i>ip- address</i> | Network that BGP advertises. |
| <i>/ prefix-length</i> | Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| <i>ip-address mask</i> | Network mask applied to the <i>ip-address</i> argument. |
| route-policy <i>route-policy-name</i> | (Optional) Specifies a route policy to use to modify the attributes of the network. |

Command Default

No networks are specified.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A network specified with this command is originated and advertised to neighbors only if there exists a route for the network in the routing table. That is, there must be a route learned using local or connected networks, static routing, or a dynamic IGP such as IS-IS or OSPF.

Other than the available system resources on the router, no limit exists on the number of network commands that can be configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the local router to originate the IPv4 unicast network 172.20.0.0/16:

```
RP/0/RSP0/CPU0:router(config)#router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# network 172.20.0.0/16
```

Related Commands

| Command | Description |
|--|---|
| network backdoor , on page 198 | Specifies a backdoor route to a BGP border router that provides better information about the network. |
| redistribute (BGP) , on page 231 | Redistributes routes from one routing domain into another routing domain. |

network backdoor

To set the administrative distance on an external Border Gateway Protocol (eBGP) route to that of a locally sourced BGP route, causing it to be less preferred than an Interior Gateway Protocol (IGP) route, use the **network backdoor** command in an appropriate configuration mode. To disable setting the administrative distance to the value for locally sourced BGP routes, use the **no** form of this command.

network *{ip-address/prefix-length | ip-address mask}* **backdoor**

no network *{ip-address/prefix-length | ip-address mask}* **backdoor**

Syntax Description

| | |
|------------------------|--|
| <i>ip-address</i> | Network that provides a backdoor route. |
| <i>/ prefix-length</i> | Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| <i>mask</i> | Network mask applied to the <i>ip-address</i> argument. |

Command Default

No backdoor routes are installed.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configuring the **network backdoor** command does not cause BGP to originate a network, even if an IGP route for the network exists. Ordinarily, the backdoor network would be learned through both an eBGP and IGP. The BGP best-path selection algorithm does not change when a network is configured as a backdoor network.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows IP Version 4 (IPv4) unicast network 192.168.40.0/24 configured as a backdoor network:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# network 192.168.40.0/24 backdoor
```

Related Commands

| Command | Description |
|---|---|
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

next-hop-self

To disable next-hop calculation and insert your own address in the next-hop field of Border Gateway Protocol (BGP) updates, use the **next-hop-self** command in an appropriate configuration mode. To enable next-hop calculation, use the **no** form of this command.

next-hop-self [**inheritance-disable**]

no next-hop-self [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | (Optional) Allows a next-hop calculation override when this feature may be inherited from a neighbor group or address family group. |
|----------------------------|---|

Command Default

When this command is not specified, the software calculates the next hop for BGP updates accepted by the router.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration
 IPv4 labeled-unicast address family configuration
 IPv6 labeled-unicast address family configuration
 VRF labeled-unicast address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

| Release | Modification |
|-------------|--|
| Release 4.0 | <p>This command was supported in the following address family configuration modes:</p> <ul style="list-style-type: none"> • IPv4 labeled-unicast • IPv6 labeled-unicast • VRF labeled-unicast |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **next-hop-self** command to set the BGP next-hop attribute of routes being advertised over a peering session to the local source address of the session.

This command is useful in nonmeshed networks in which BGP neighbors may not have direct access to all other neighbors on the same IP subnet.

If this command is configured for a neighbor group or address family group, a neighbor using the group inherits the configuration. Configuring the command specifically for a neighbor overrides any inherited value.

Configuring the **next-hop-self** command under IPv4 labeled-unicast, IPv6 labeled-unicast, or VRF labeled-unicast address family configuration mode enables next-hop-self for labeled prefixes advertised to an iBGP peer.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the next hop of the update field for all IP Version 4 (IPv4) unicast routes advertised to neighbor 172.20.1.1 to an address of the local router:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# next-hop-self
```

The following example shows how to disable the **next-hop-self** command for neighbor 172.20.1.1. If not overridden, the next hop would be inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# next-hop-self
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
```

```
RP/0/RSP0/CPU0:router(config-bgp-nbr-af) # next-hop-self inheritance-disable
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| use, on page 514 | Inherits characteristics from a neighbor group, session group, or address family group. |

next-hop-unchanged

To disable overwriting of the next hop before advertising to external Border Gateway Protocol (eBGP) peers, use the **next-hop-unchanged** command in an appropriate configuration mode. To enable overwriting of the next hop, use the **no** form of this command.

next-hop-unchanged [**inheritance-disable**]

no next-hop-unchanged [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|--|
| inheritance-disable | (Optional) Allows overwriting of the next hop before advertising to eBGP peers when this feature may be inherited from a neighbor group or address family group. |
|----------------------------|--|

Command Default

Overwriting of the next hop is allowed.

Command Modes

VPNv4 address family group configuration
 VPNv4 neighbor address family configuration
 VPNv4 neighbor group address family configuration
 IPv4 labeled-unicast address family configuration
 IPv6 labeled-unicast address family configuration
 IPv4 address family configuration
 IPv6 address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |
| Release 4.0.0 | This command was supported in the following address family configuration modes: <ul style="list-style-type: none"> • IPv4 labeled-unicast address family configuration • IPv6 labeled-unicast address family configuration • IPv4 unicast address family configuration • IPv6 unicast address family configuration |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **next-hop-unchanged** command to propagate the next hop unchanged for multihop eBGP peering sessions. This command should not be configured on a route reflector, and the **next-hop-self** command should not be used to modify the next-hop attribute for a route reflector when this feature is enabled for a route reflector client.

**Note**

Incorrectly setting BGP attributes for a route reflector can cause inconsistent routing, routing loops, or a loss of connectivity. Setting BGP attributes for a route reflector should be attempted only by an experienced network operator.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to disable the overwriting of next hops before advertising to eBGP peers:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# next-hop-unchanged disable
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
```

Related Commands

| Command | Description |
|---|---|
| next-hop-self , on page 200 | Disables next-hop calculation and allows you to insert your own address in the next-hop field of BGP updates. |
| use , on page 514 | Inherits characteristics from a neighbor group, session group, or address family group. |

nexthop resolution prefix-length minimum

To set minimum prefix-length for nexthop resolution, use the **nexthop resolution prefix-length minimum** command in an appropriate configuration mode. To disable the minimum prefix-length for nexthop resolution, use the **no** form of this command.

nexthop resolution prefix-length minimum *prefix-length-value*

no nexthop resolution prefix-length minimum *prefix-length-value*

Syntax Description

| | |
|----------------------------|---|
| <i>prefix-length-value</i> | Sets the minimum prefix-length. Range is 0 to 32. |
|----------------------------|---|

Command Default

Nexthop resolution for minimum prefix-length is disabled.

Command Modes

VPNv4 Unicast address family
VRF IPv4 Unicast address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the minimum prefix-length for nexthop resolution as 32:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#nexthop resolution prefix-length minimum 32
```

nexthop route-policy

To specify that BGP routes are resolved using only next hops whose routes match specific characteristics, use the **nexthop route-policy** command in the appropriate configuration mode. To remove the **nexthop route-policy** command from the configuration file and restore the system to its default behavior, use the **no** form of this command.

nexthop route-policy *route-policy-name*

no nexthop route-policy *route-policy-name*

Syntax Description

| | |
|--------------------------|---|
| <i>route-policy-name</i> | Route policy to use for filtering based on next hops. |
|--------------------------|---|

Command Default

No default behavior or values

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VPNv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nexthop route-policy** command to configure route policy filtering using next hops.

The BGP next-hop tracking feature allows you to specify that BGP routes are resolved using only next hops whose routes have the following characteristics:

- To avoid the aggregate routes, the prefix length must be greater than a specified value.
- The source protocol must be from a selected list, ensuring that BGP routes are not used to resolve next hops that could lead to oscillation.

This route policy filtering is possible because RIB identifies the source protocol of a route that resolves a next hop as well as the mask length associated with the route.

The next-hop attach point supports matching using the protocol name and mask length. BGP marks all next hops that are rejected by the route policy as invalid, and no best path is calculated for the routes that use the invalid next hop. The invalid next hops continue to stay in the active cache and can be displayed as part of the **show bgp nexthop** command with an invalid status.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to specify the route policy nexthop_A as the policy to use for filtering next hops:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# nexthop route-policy nexthop_A
```

Related Commands

| Command | Description |
|--|--|
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |
| show bgp nexthops, on page 377 | Display statistical information about the BGP next hops. |

nexthop trigger-delay

To specify the delay for triggering next-hop calculations, use the **nexthop trigger-delay** command in the appropriate configuration mode. To set the trigger delay to the default value, use the **no** form of this command.

nexthop trigger-delay {*critical delay*| *non-critical delay*}

no nexthop trigger-delay {*critical delay*| *non-critical delay*}

Syntax Description

| | |
|---------------------|---|
| critical | Specifies critical next-hop events. For example, when the next hop is unreachable. |
| <i>delay</i> | Trigger delay, in milliseconds. Range is 0 to 4294967295. |
| non-critical | Specifies noncritical next-hop events. For example, Interior Gateway Protocol (IGP) metric changes. |

Command Default

critical : 3000 msec for IPv4 address family and IPv6 address family

non-critical : 10000 msec IPv4, IPv6, VPNv4, and VPNv6 address families

Command Modes

IPv4 address family configuration

Pv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nexthop trigger-delay** command to allow for a dynamic way for Interior Gateway Protocol (IGP) to converge. This convergence allows BGP to accumulate all notifications and trigger fewer walks, resulting in fewer interprocess communications (IPCs) to the Routing Information Base (RIB) for route addition, deletion, and modification and fewer updates to peers.



Note

A high *delay* value can be configured to effectively turn off next-hop tracking.

The **non-critical** *delay* value must always be set to at least equal or greater than the **critical** *delay value*.

The *delay* should be slightly higher than the time it takes for the IGP to settle into a steady state after some event (IGP convergence time).

Task ID

| Task ID | Operations |
|----------------|-------------------|
| bgp | read, write |

Examples

The following example shows how to set the critical next-hop trigger delay to 3500 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109  
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-af)# nexthop trigger-delay critical 3500
```

nsr (BGP)

To activate Border Gateway Protocol (BGP) nonstop routing (NSR), use the **nsr** command in BGP global configuration mode. To deactivate BGP NSR, use the **no** form of this command.

nsr

no nsr

Syntax Description This command has no arguments or keywords.

Command Default BGP NSR is not activated.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|-----------------------------------|
| | Release 3.9.0 | This command was introduced. |
| | Release 4.1.0 | 5000 NSR sessions were supported. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nsr** command to enable the Border Gateway Protocol (BGP) Nonstop Routing (NSR) with Stateful Switchover (SSO). This enables all **bgp** peerings to maintain the BGP state to ensure continuous packet forwarding during events that could interrupt service.

BGP supports 5000 NSR sessions.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to enable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 120
RP/0/RSP0/CPU0:router(config-bgp)# nsr
```

The following example shows how to disable BGP NSR:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router bgp 120  
RP/0/RSP0/CPU0:router(config-bgp)# no nsr
```

Related Commands

| Command | Description |
|---|--|
| router bgp, on page 253 | Configures the Border Gateway Protocol (BGP) routing process. |
| nsr process-failures switchover | Configures failover as a recovery action in case of process failures for active instances to switch over to a standby route processor (RP) or a standby distributed route processor (DRP) to maintain nonstop routing (NSR). |
| show bgp nsr, on page 386 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

orf

To specify Outbound Route Filter (ORF) and inbound filtering criteria, use the **orf route-policy** command in an appropriate configuration mode. To restore the system to its default condition, use the **no** form of this command.

orf route-policy *route-policy-name*

no orf route-policy *route-policy-name*

Syntax Description

| | |
|--------------------------|---------------------------|
| <i>route-policy-name</i> | Name of the route policy. |
|--------------------------|---------------------------|

Command Default

No ORF route policy is defined.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 IPv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure outbound and inbound filtering criteria:

```
RP/0/RSP0/CPU0:router(config)#router bgp 6  
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)#orf route-policy policy_A
```

Related Commands

| Command | Description |
|---|--|
| route-policy (BGP), on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |

password (BGP)

To enable Message Digest 5 (MD5) authentication on a TCP connection between two Border Gateway Protocol (BGP) neighbors, use the **password** command in an appropriate configuration mode. To disable MD5 authentication, use the **no** form of this command.

password {**clear**|**encrypted**} *password*

no password [**clear** *password*|**encrypted** *password*]

Syntax Description

| | |
|------------------|--|
| clear | Specifies that an unencrypted password follows. The password must be a case-sensitive, clear-text unencrypted password. |
| encrypted | Specifies that an encrypted password follows. The password must be a case-sensitive, encrypted password. |
| <i>password</i> | Password of up to 80 characters. The password can contain any alphanumeric characters. However, if the first character is a number or the password contains a space, the password must be enclosed in double quotation marks; for example, "2 password." |

Command Default

When this command is not specified in the appropriate configuration mode, MD5 authentication is not enabled on a TCP connection between two BGP neighbors.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configure a password to enable authentication between two BGP peers. Use the **password** command to verify each segment sent on the TCP connection between the peers. The same password must be configured on both networking devices, otherwise a connection cannot be made. The authentication feature uses the MD5

algorithm. Specifying this command causes the software to generate and check the MD5 digest on every segment sent on the TCP connection.

Configuring a neighbor password does not cause the existing session for a neighbor to end. However, until the new password is configured on the remote router, the local BGP process does not receive keepalive messages from the remote device. If the password is not updated on the remote device by the end of the hold time, the session ends. The hold time can be changed using the **timers** command or the **timers bgp** command.

If this command is configured for a neighbor group or neighbor address family group, a neighbor using the group inherits the configuration. Values of commands configured specifically for a neighbor overrides inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure neighbor 172.20.1.1 to use MD5 authentication with the password password1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)#neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#password clear password1
```

Related Commands

| Command | Description |
|--|--|
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| password-disable , on page 218 | Overrides any inherited password configuration from a neighbor group or session group for BGP neighbors. |
| session-group , on page 264 | Creates a session group and enters session group configuration mode. |
| timers (BGP) , on page 491 | Set the timers for a specific BGP neighbor. |

password (rpki-server)

To specify a SSH password for the RPKI cache-server, use the **password** command in rpki-server configuration mode. To remove the SSH passwords, use the **no** form of this command.

password *password*

no password *password*

Syntax Description

| | |
|-----------------|---|
| <i>password</i> | Enters a password to be used for the SSH transport mechanism. |
|-----------------|---|

Command Default

Password is not configured.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

SSH expects to use an authentication method to connect to a remote server. The SSH authentication method to connect to RPKI server is password-based. So, the RPKI cache-server must be configured with username and password. A username and password must be configure for each server configured under BGP that uses the SSH transport

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure a username (*rpki-user*) and password (*rpki-ssh-pass*) for the RPKI cache-server SSH transport mechanism:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
```

```
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
```

password-disable

To override any inherited password configuration from a neighbor group or session group for Border Gateway Protocol (BGP) neighbors, use the **password-disable** command in an appropriate configuration mode. To disable overriding any inherited password command, use the **no** form of this command.

password-disable

no password-disable

Syntax Description This command has no arguments or keywords.

Command Default Configured passwords for neighbor and session groups are inherited.

Command Modes

- Neighbor configuration
- VRF neighbor configuration
- Neighbor group configuration
- Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you specify a password on a neighbor group or session group, all users of the group inherit the password. Specifying a different **password** command specifically on a neighbor that uses the group overrides the inherited value. Specifying **password-disable** on a neighbor that uses the group disables password authentication for the neighbor.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to disable MD5 authentication for neighbor 172.20.1.1, preventing it from inheriting the password password1 from session group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
```

```

RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# password clear password1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# password-disable

```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| password (BGP), on page 214 | Enables MD5 authentication on a TCP connection between two BGP neighbors. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| use, on page 514 | Inherits characteristics from a neighbor group, a session group, or an address family group. |

permanent-network

To define a prefix set as permanent, use the **permanent-network** command in the global address family configuration mode. To remove a prefix set as permanent, use the **no** form of this command. The **permanent-network** command uses a route-policy to identify the set of prefixes (networks) for which permanent paths needs to be created.

The permanent network feature supports only prefixes in IPv4 unicast and IPv6 unicast address-families under the default Virtual Routing and Forwarding (VRF).

permanent-network route-policy route-policy-name

no permanent-network

Syntax Description

| | |
|---------------------------------------|--|
| route-policy route-policy-name | Specifies a configured routing policy. |
|---------------------------------------|--|

Command Default

None

Command Modes

Address-family configuration.

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 5.1.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to define permanent path for a route policy named POLICY-PERMANENT-NETWORK-IPv4:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-af)# permanent-network route-policy POLICY-PERMANENT-NETWORK-IPv4
```


precedence

To set the precedence level, use the **precedence** command in the appropriate configuration mode. To remove the **precedence** command from the configuration file and restore the system to its default interval values, use the **no** form of this command.

precedence *value*

no precedence [*value*]

Syntax Description

value Value of the precedence. The precedence value can be a number from 0 to 7, or it can be one of the following keywords:

critical —Set packets with critical precedence (5)

flash — Set packets with flash precedence (3)

flash-override —Set packets with flash override precedence (4)

immediate —Set packets with immediate precedence (2)

internet —Set packets with internetwork control precedence (6)

network —Set packets with network control precedence (7)

priority —Set packets with priority precedence (1)

routine —Set packets with routine precedence (0)

Command Default

No default behavior or values

Command Modes

Neighbor configuration

Neighbor session group configuration

Neighbor group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **precedence** command to set the precedence value.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to set the precedence to 2:

```
RP/0/RSP0/CPU0:router(config)# router bgp 5  
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.1.1.1  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100  
RP/0/RSP0/CPU0:router(config-bgp-nbr)# precedence 2
```

preference (rpki-server)

To specify a preference value for the RPKI cache-server, use the **preference** command rpki-server configuration mode. To remove the preference value, use the **no** form of this command.

preference *preference-value*

no preference *preference-value*

Syntax Description

preference-value

Specifies a RPKI cache preference value. Range is 1 to 10.

Note A lower value is recommended

Command Default

Preference value is not set.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set preference value for RPKI configuration as 1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router (config)#router bgp 100
RP/0/RSP0/CPU0:router (config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router (config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router (config-bgp-rpki-cache)#username rpki-user
RP/0/RSP0/CPU0:router (config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RSP0/CPU0:router (config-bgp-rpki-cache)#preference 1
```

purge-time (rpki-server)

To configure the time BGP waits to keep routes from RPKI cache-server after the cache session drops, use the **purge-time** command in rpki-server configuration mode. To remove the purge-time configuration, use the **no** form of this command.

purge-time *time-in-seconds*

no purge-time *time-in-seconds*

Syntax Description

| | |
|------------------------|---|
| <i>time-in-seconds</i> | Sets the purge time in seconds. Range is 30 to 360 seconds. |
|------------------------|---|

Command Default

Purge time is not set.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a cache session is dropped then a "purge-timer" is started for that cache. If the session re-establishes within the timer interval, then the purge timer is stopped and no further action is taken. If the cache session does not re-establish within the timer interval, only then does BGP remove all ROAs from the cache.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the purge-time for RPKI cache as 30 seconds:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki cache 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
```

```
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#preference 1  
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
```

rd

To configure a route distinguisher, use the **rd** command in VRF configuration mode. To disable the route distinguisher, use the **no** form of this command.

rd {*as-number* : *nn*| *ip-address* : *nn*| **auto**}

no rd {*as-number* : *nn*| *ip-address* : *nn*| **auto**}

Syntax Description

| | |
|----------------------|--|
| <i>as-number:nn</i> | <ul style="list-style-type: none"> • <i>as-number</i>—16-bit Autonomous system (AS) number of the route distinguisher <ul style="list-style-type: none"> ◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. • <i>nn</i>—32-bit number |
| <i>ip-address:nn</i> | <p>IP address of the route distinguisher.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—32-bit IP address • <i>nn</i>—16-bit number |
| auto | Automatically assigns a unique route distinguisher. |

Command Default

No default behavior or values

Command Modes

VRF configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rd** command to make the prefix unique across multiple VRFs.

Auto assignment of route distinguishers can be done only if a router ID is assigned using the **bgp router-id** command in BGP router configuration mode. The unique router ID is used for automatic route distinguisher generation.

The following are restrictions when configuring route distinguishers:

- BGP router-id must be configured before **rd auto** can be configured
- Route distinguisher cannot be changed or removed when an IPv4 unicast address family is configured under VRF.
- BGP router-id cannot be changed or removed when **rd auto** is configured under a VRF.
- When **rd auto** is configured under a VRF, the IP address for the router distinguisher configured under another VRF must be different from that of the BGP router-id
- If a route distinguisher with same IP address as BGP router-id exists, the **rd auto** is not permitted.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to automatically assign a unique route distinguisher to VRF instance vrf-1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-bgp-vrf)# rd auto
```

Related Commands

| Command | Description |
|--|--|
| bgp router-id, on page 101 | Configures a fixed router ID for a BGP-speaking router. |
| export route-target, on page 163 | Configures a VRF export route-target extended community. |
| import route-target, on page 167 | Configures a VRF import route-target extended community. |

receive-buffer-size

To set the size of the receive buffers for a Border Gateway Protocol (BGP) neighbor, use the **receive-buffer-size** command in an appropriate configuration mode. To remove the **receive-buffer-size** command from the configuration file and restore the system to its default condition in which the software uses the default size, use the **no** form of this command.

receive-buffer-size *socket-size* [*bgp-size*]

no receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

| | |
|--------------------|--|
| <i>socket-size</i> | Size, in bytes, of the receive-side socket buffer. Range is 512 to 131072. |
| <i>bgp-size</i> | (Optional) Size, in bytes, of the receive buffer in BGP. Range is 512 to 131072. |

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because it allows the software to process a larger number of packets simultaneously. However, allocating larger buffers consumes more memory on the router.



Note

Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the receive buffer sizes for neighbor 172.20.1.1 to be 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# receive-buffer-size 65536 8192
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-buffer-size, on page 258 | Sets the size of the send buffers for a BGP neighbor. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| socket receive-buffer-size, on page 479 | Sets the size of the receive buffers for all BGP neighbors. |

redistribute (BGP)

To redistribute routes from one routing domain into Border Gateway Protocol (BGP), use the **redistribute** command in an appropriate configuration mode. To disable route redistribution, use the **no** form of this command.

Connected

redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute connected [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Enhanced Interior Gateway Routing Protocol (EIGRP)

redistribute eigrp *process-id* [**match** {**external**|**internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute eigrp *process-id* [**match** {**external**|**internal**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Intermediate System-to-Intermediate System (IS-IS)

redistribute isis *process-id* [**level** {**1**|**1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute isis *process-id* [**level** {**1**|**1-inter-area** | **2**}] [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Open Shortest Path First (OSPF)

redistribute ospf *process-id*

no redistribute ospf *process-id*

Routing Information Protocol

redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute rip [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Static

redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]

no redistribute static [**metric** *metric-value*] [**route-policy** *route-policy-name*]

Syntax Description

| | |
|-----------------------------------|--|
| connected | Redistributes connected routes. Connected routes are established automatically when IP is enabled on an interface. |
| metric <i>metric-value</i> | (Optional) Specifies the Multi Exit Discriminator (MED) attribute used for the redistributed route. Range is 0 to 4294967295. Use a value consistent with the destination protocol. By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0. |

| | |
|---|---|
| route-policy <i>route-policy-name</i> | (Optional) Specifies a configured routing policy to filter redistributed routes. A route policy is used to filter the importation of routes from this source routing protocol to BGP. |
| eigrp | Specifies that routes are distributed from EIGRP. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode. |
| <i>process-id</i> | <p>For the eigrp keyword, an EIGRP instance name from which routes are to be redistributed.</p> <p>For the isis keyword, an IS-IS instance name from which routes are to be redistributed.</p> <p>For the ospf keyword, an OSPF instance name from which routes are to be redistributed.</p> <p>The <i>process-id</i> value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| match { internal external [1 2] nssa-external [1 2] } | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> • internal —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • nssa-external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> |
| isis | <p>Specifies that routes are distributed from the IS-IS protocol.</p> <p>Redistribution from IS-IS is allowed under IPv4 unicast, IPv4 multicast, IPv6 unicast, and address-families. Redistribution is not allowed under VPNv4 address-families.</p> |
| level { 1 1-inter-area 2 } | <p>(Optional) Specifies the IS-IS level from which routes are redistributed. It can be one of the following:</p> <ul style="list-style-type: none"> • 1 —Routes are redistributed from Level 1 routes. • 1-inter-area —Routes are redistributed from Level 1 interarea routes. • 2 —Routes are redistributed from Level 2 routes. |
| ospf | Specifies that routes are distributed from the OSPF protocol. You must be in IPv4 unicast or multicast address family configuration mode or in VRF IPv4 address family configuration mode. |

| | |
|---------------|---|
| rip | Specifies that routes are distributed from RIP. You must be in IPv4 unicast or multicast address family configuration mode. |
| static | Redistributes IP static routes. |

Command Default

Route redistribution is disabled.

For IS-IS, the default is to redistribute Level 1 and Level 2 routes.

For OSPF, the default is to redistribute internal, external, and NSSA external routes of Type 1 and Type 2.

By default, the Interior Gateway Protocol (IGP) metric is assigned to the route. For connected and static routes the default metric is 0.

metric *metric-value*: 0

match { **internal** | **external** [1 | 2] | **nssa-external** [1 | 2]}: If no match is specified, the default is to match all routes.

Command Modes

IPv4 address family configuration, both unicast and multicast (**connected**, **eigrp**, **isis**, **ospf**, **rip**, and **static** are supported)

IPv6 address family configuration, both unicast and multicast (**connected**, **eigrp**, **isis**, **ospfv3**, and **static** are supported)

VRF IPv4 address family configuration (**connected** , **eigrp** , **ospf** , **rip** , and **static** are supported)

VRF IPv6 address family configuration (**connected** , **eigrp** , and **static** are supported)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

When redistributing routes (into BGP) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Each instance of a protocol may be redistributed independently of the others. Changing or removing redistribution for a particular instance does not affect the redistribution capability of other protocols or other instances of the same protocol.

Networks specified using the **network** command are not affected by the **redistribute** command; that is, the routing policy specified in the **network** command takes precedence over the policy specified through the **redistribute** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to redistribute IP Version 4 (IPv4) unicast OSPF routes from OSPF instance 110 into BGP:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# redistribute ospf 110
```

Related Commands

| Command | Description |
|---|---|
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |

refresh-time (rpki-server)

To configure the time BGP waits in between sending periodic serial queries to the RPKI server, use the **refresh-time** command in rpki-server configuration mode. To remove the refresh-time configuration, use the **no** form of this command.

refresh-time {*time-in-seconds*| **off**}

no refresh-time {*time-in-seconds*| **off**}

Syntax Description

| | |
|------------------------|--|
| off | Specifies not to send serial queries periodically. |
| <i>time-in-seconds</i> | Sets the refresh-time in seconds. Range is 30 to 3600 seconds. |

Command Default

Refresh-time is not set.

Command Modes

RPKI cache configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the refresh-time for BGP to wait in between sending periodic serial queries to the server as 30 seconds:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#preference 1
```

```
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30  
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#refresh-time 30
```

response-time (rpki-server)

To configure the time BGP waits for a response from the RPKI cache-server after sending a serial or reset query, use the **response-time** command in rpki-server configuration mode. To remove the response-time configuration, use the **no** form of this command.

response-time {*time-in-seconds*| **off**}

no response-time {*time-in-seconds*| **off**}

Syntax Description

| | |
|------------------------|--|
| off | Specifies to wait indefinitely for a response from the RPKI cache. |
| <i>time-in-seconds</i> | Specifies the response-time in seconds. Range is 30 to 3600 seconds. |

Command Default

Response-time is not set.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the time for BGP to wait for a response from the RPKI server as 30 seconds, after sending a serial or reset query:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 72.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#preference 1
```

response-time (rpki-server)

```
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30  
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#refresh-time 30  
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#response-time 30
```

remote-as (BGP)

To create a Border Gateway Protocol (BGP) neighbor and begin the exchange of routing information, use the **remote-as** command in an appropriate configuration mode. To delete the entry for the BGP neighbor, use the **no** form of this command.

remote-as *as-number*

no remote-as [*as-number*]

Syntax Description

| | |
|------------------|---|
| <i>as-number</i> | Autonomous system (AS) to which the neighbor belongs. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
|------------------|---|

Command Default

No BGP neighbors exist.

Command Modes

Neighbor configuration
VRF neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **remote-as** command to create a neighbor and assign it a remote autonomous system number. A neighbor must have a remote autonomous system number before any other commands can be configured for it. Removing the remote autonomous system from a neighbor causes the neighbor to be deleted. You cannot remove the autonomous system number if the neighbor has other configuration.

**Note**

We recommend that you use the **no neighbor** command rather than the **no remote-as** command to delete a neighbor.

A neighbor specified with a remote autonomous system number that matches the autonomous system number specified in the **router bgp** command identifies the neighbor as internal to the local autonomous system. Otherwise, the neighbor is considered external.

Configuration of the **remote-as** command for a neighbor group or session group using the **neighbor-group** command or **session-group** command causes all neighbors using the group to inherit the characteristics configured with the command. Configuring the command directly for the neighbor overrides the value inherited from the group.

In the neighbor configuration submode, configuring use of a session group or neighbor group for which **remote-as** is configured creates a neighbor and assigns it an autonomous system number if the neighbor has not already been created.

**Note**

Do not combine **remote-as** commands and **no use neighbor-group** commands, or **remote-as** commands and **no use session-group** commands, in the same configuration commit.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to assign autonomous system numbers on two neighbors, neighbor 10.0.0.1, (internal) and neighbor 192.168.0.1 (external), setting up a peering session that shares routing information between this router and each of these neighbors:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group group2
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)#exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)#use session-group group2
```

The following example shows how to configure a session group called group2 with an autonomous system number 1. Neighbor 10.0.0.1 is created when it inherits the autonomous system number 1 from session group group2.

```
RP/0/RSP0/CPU0:router(config)#router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group group2
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group2
```

Related Commands

| Command | Description |
|--|---|
| neighbor (BGP) , on page 192 | Enters neighbor configuration mode for configuring BGP routing sessions. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| router bgp , on page 253 | Configures the BGP routing process. |
| session-group , on page 264 | Creates a session group and enters session group configuration mode. |
| use , on page 514 | Inherits characteristics from a neighbor group, session group, or address family group. |

remove-private-as

To remove private autonomous system numbers from autonomous system paths when generating updates to external neighbors, use the **remove-private-as** command in an appropriate configuration mode. To place the router in the default state in which it does not remove private autonomous system numbers, use the **no** form of this command.

remove-private-as [**inheritance-disable**] [**entire-aspath**]

no remove-private-as [**inheritance-disable**] [**entire-aspath**]

Syntax Description

| | |
|----------------------------|--|
| inheritance-disable | (Optional) Permits the feature to be disabled from a neighbor group or address family group instead of being inherited. |
| entire-aspath | (Optional) Removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private. |

Command Default

When this command is not specified in the appropriate configuration mode, private autonomous system numbers are not removed from updates sent to external neighbors.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |
| Release 3.9.2 | The entire-aspath keyword was supported. |

| Release | Modification |
|---------------|---|
| Release 4.1.0 | This command was supported on ASR 9000 Ethernet Line Card (Cisco ASR 9000's A9K-SIP-700). |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This feature is available for external BGP (eBGP) neighbors only.

When an update is passed to the external neighbor, the software drops any leading autonomous system sequence in the autonomous system path if the sequence contains only private autonomous system numbers and does not contain the autonomous system number of the neighbor.

If this command is used in a BGP confederation, the element following the confederation portion of the autonomous system path, if a sequence, is considered the leading sequence.

The private autonomous system values range from 64512 to 65535.

If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Use the **entire-aspath** to removes the entire private autonomous system numbers from an autonomous system path only if all ASes in the path are private.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows a configuration that removes the private autonomous system number from the IP Version 4 (IPv4) unicast updates sent to 172.20.1.1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# remove-private-as
```

The following example shows how to disable the remove private autonomous system number feature for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# remove-private-as
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# remove-private-as inheritance-disable
```

Related Commands

| Command | Description |
|---|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| remote-as (BGP) , on page 239 | Allows entries to the BGP neighbor table. |

retain local-label

To retain the local label until the network is converged, use the **retain local-label** command in an appropriate address family configuration mode. To disable the retaining of the local label, use the **no** form of this command.

retain local-label *minutes*

no retain local-label

Syntax Description

| | |
|----------------|---|
| <i>minutes</i> | Local retention time in minutes. The range is 3 to 60 minutes. The default retention time is 5 minutes. |
|----------------|---|

Command Default

minutes : 5

Command Modes

L2VPN address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable local label retention for 5 minutes:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# retain local-label 5
```

Related Commands

| Command | Description |
|---|--|
| additional-paths install backup , on page 9 | Installs a backup path into the forwarding table |
| advertise best-external , on page 21 | Advertises the best-external path to the iBGP and route-reflector peers. |

retain route-target

To accept received updates with specified route targets, use the **retain route-target** command in an appropriate configuration mode. To disable the retaining of routes tagged with specified route targets, use the **no** form of this command.

retain route-target {**all**| **route-policy** *route-policy-name*}

no retain route-target [**all**| **route-policy** *route-policy-name*]

Syntax Description

| | |
|---|---|
| all | Accepts received updates containing at least one route target. |
| route-policy <i>router-policy-name</i> | Accepts received updates accepted by a specified route filter policy. |

Command Default

The default is to accept all route targets.

Command Modes

VPNv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **retain route-target** command to configure a route reflector (RR) to retain routes tagged with specific route targets (RT).

The **retain route-target** is a required command for Inter-AS option B ASBR. By default, an Inter-AS option B ASBR needs the **retain route-target** configured to get VPNv4 BGP table from PE routers, either with the **all** or with the **route-policy** option.

A provider edge (PE) router is not required to hold all VPNv4 routes. The PE router holds only routes that match the import RT of the VPNs configured on it, but a RR must retain all VPNv4 routes because it may peer with PE routers and different PEs may require different RT-tagged VPNv4 routes. Configuring an RR to hold only routes that have a defined set of RT communities and configuring some of these RRs to service a different set of VPNs provides scalability to the RRs. A PE can be configured to peer with all RRs that service the VPN routing and forwarding (VRF) instances configured on the PE. When a new VRF is configured with an RT for which the PE does not already hold routes, the PE issues route refresh requests to the RRs and gets the relevant VPN routes.

The **route-policy** *route-policy-name* keyword and argument takes the policy name that lists the extended communities that a path should have for the RR to retain the path.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure RR to retain all routes with the route filter policy ft-policy-A:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)# retain route-target route-filter ft-policy-A
```

Related Commands

| Command | Description |
|--|--|
| import route-target, on page 167 | Configures a VRF import route-target extended community. |

route-policy (BGP)

To apply a routing policy to updates advertised to or received from a Border Gateway Protocol (BGP) neighbor, use the **route-policy** command in an appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* [*parameter1*, *parameter2*, . . . , *parameterN*] {**in** | **out**}

no route-policy *route-policy-name* [*parameter1*, *parameter2*, . . . , *parameterN*] {**in** | **out**}

Syntax Description

| | |
|--------------------------|---|
| <i>route-policy-name</i> | Name of route policy. Up to 16 parameters can follow the route-policy-name, enclosed in brackets ([]). |
| in | Applies policy to inbound routes. |
| out | Applies policy to outbound routes. |

Command Default

No policy is applied.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes. The **route-policy** command is used to define a policy.



Note Configuring a large number of uniquely named outbound neighbor policies can adversely affect performance. This is true even if the uniquely named route policies are functionally identical. The user is discouraged from configuring multiple functionally identical route policies for use with this command. For example, if Policy A and Policy B are identical but named for different neighbors, the two policies should be configured as a single policy.

If the **route-policy** command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to apply the In-Ipv4 policy to inbound IP Version 4 (IPv4) unicast routes from neighbor 172.20.1.1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-policy In-Ipv4 in
```

Related Commands

| Command | Description |
|--|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

route-reflector-client

To configure the router as a Border Gateway Protocol (BGP) route reflector and configure the specified neighbor as its client, use the **route-reflector-client** command in an appropriate configuration mode. To disable configuring the neighbor as a client, use the **no** form of this command.

route-reflector-client [**inheritance-disable**]

no route-reflector-client [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | (Optional) Allows the configuration inherited from a neighbor group or address family group to be overridden. |
|----------------------------|---|

Command Default

The neighbor is not treated as a route reflector client.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is restricted to internal BGP (iBGP) neighbors only.

Use the **route-reflector-client** command to configure the local router as the route reflector and the specified neighbor as one of its clients. All neighbors configured with this command are members of the client group, and the remaining iBGP peers are members of the nonclient group for the local route reflector.

By default, all iBGP speakers in an autonomous system must be fully meshed with each other, and neighbors do not readvertise iBGP learned routes to other iBGP neighbors.

With route reflection, all iBGP speakers need not be fully meshed. An iBGP speaker, the route reflector, passes learned iBGP routes to some number of iBGP client neighbors. Learned iBGP routes eliminate the need for each router running BGP to communicate with every other device running BGP in the autonomous system.

The local router is a route reflector as long as it has at least one route reflector client.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows neighbor at 172.20.1.1 configured as a route reflector client for IP Version 4 (IPv4) unicast routes:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client
```

The following example disables the route-reflector client for neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# route-reflector-client
RP/0/RSP0/CPU0:router(config-bgp-afgrp)#exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 140
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# route-reflector-client inheritance-disable
```

Related Commands

| Command | Description |
|--|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| bgp cluster-id , on page 70 | Configures the cluster ID if the BGP cluster has more than one route reflector. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in global configuration mode. To remove all BGP configurations and terminate the BGP routing process, use the **no** form of this command.

router bgp *as-number* [**instance** *instance-name*]

no router bgp [*as-number*]

Syntax Description

| | |
|---|--|
| <i>as-number</i> | Number that identifies the autonomous system (AS) in which the router resides. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| instance <i>instance-name</i> | Specifies an instance and instance name. The maximum length for the instance name is 32 characters. The router bgp instance <i>instance-name</i> command replaced the distributed speaker command. |

Command Default

No BGP routing process is enabled.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system number notation was supported. |
| Release 4.2.0 | The instance and <i>instance-name</i> keyword and argument were added to support BGP Multi-Instance/Multi-AS feature. The command with the instance and <i>instance-name</i> keyword and argument replaced the distributed speaker command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **router bgp** command to set up a distributed routing core that automatically guarantees the loop-free exchange of routing information between autonomous systems.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| rib | read, write |

Examples

The following example shows how to configure a BGP process for autonomous system 120:

```
RP/0/RSP0/CPU0:router(config)# router bgp 120
```

rpki server

To enter resource public key infrastructure (RPKI) cache-server (rpki-sever) configuration mode and enable rpki parameters configuration, use the **rpki server** command in Router BGP configuration mode. To remove the rpki-server configuration mode and delink cache-server from the cache list, use the **no** form of this command.

rpki server {*host-name*| *ip-address*}

no rpki server {*host-name*| *ip-address*}

Syntax Description

| | |
|-------------------|---------------------------------------|
| <i>host-name</i> | Host name of the RPKI cache database. |
| <i>ip-address</i> | IP Address of the RPKI cache databse. |

Command Default

RPKI server configuration is disabled.

Command Modes

Router BGP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure an rpki cache-server database and enter rpki-server configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#
```

selective-vrf-download disable

To disable selective VRF download (SVD) on a line card to enable download all prefixes and labels to the line card, use the **selective-vrf-download disable** command in global configuration mode. To enable the SVD, use the **no** form of this command.

selective-vrf-download disable

no selective-vrf-download disable

Syntax Description This command has no keywords or arguments.

Command Default SVD is enabled.

Command Modes global configuration

| Command History | Release | Modification |
|-----------------|---------------|---|
| | Release 4.2.0 | This command was introduced. |
| | Release 4.3.1 | Support for this command was removed. This command was replaced with the no svd platform enable command. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must failover the active RP or reload the router after disabling SVD for the configuration change to get activated.

| Task ID | Task ID | Operation |
|---------|-------------|-------------|
| | ip-services | read, write |

Examples This example shows how to disable selective vrf download:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#selective-vrf-download disable
```

Related Commands

| Command | Description |
|---|--|
| show svd role, on page 470 | Displays Selective VRF Download (SVD) role information. |
| show svd state, on page 472 | Displays Selective VRF Download (SVD) state information. |

send-buffer-size

To set the size of the send buffers for a Border Gateway Protocol (BGP) neighbor, use the **send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default values, use the **no** form of this command.

send-buffer-size *socket-size* [*bgp-size*]

no send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

| | |
|--------------------|---|
| <i>socket-size</i> | Size, in bytes, of the send-side socket buffer. Range is 4096 to 131072. |
| <i>bgp-size</i> | (Optional) Size, in bytes, of the BGP process send buffer. Range is 4096 to 131072. |

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Use the **socket send-buffer-size** command to change the defaults.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **send-buffer-size** command to increase the buffer size employed when sending updates to a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on the router.



Note

Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses more memory indefinitely.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the send buffer sizes for neighbor 172.20.1.1 to be 8192 bytes for both the socket buffer and the BGP buffer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# send-buffer-size 8192 8192
```

Related Commands

| Command | Description |
|--|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| receive-buffer-size, on page 229 | Sets the size of the receive buffers for a BGP neighbor. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| socket send-buffer-size, on page 481 | Sets the size of the send buffers for all BGP neighbors. |

send-community-ebgp

To specify that community attributes should be sent to an external Border Gateway Protocol (eBGP) neighbor, use the **send-community-ebgp** command in an appropriate configuration mode. To disable sending community attributes to an eBGP neighbor, use the **no** form of this command.

send-community-ebgp [**inheritance-disable**]

no send-community-ebgp [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | (Optional) Allows configuration inherited from a neighbor group or address family group to be overridden. |
|----------------------------|---|

Command Default

Community (COMM) attributes are NOT sent to eBGP peers (including PE-CE peers).

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 IPv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration
 VPNv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **send-community-ebgp** command to control whether community attributes are sent to eBGP neighbors. This command cannot be configured for iBGP neighbors as community attributes are always sent to iBGP neighbors.

When IOS XR BGP updates community attributes for eBGP VPN peers (VPNv4 or VPNv6), there is no need to configure the **send-community-ebgp** command separately. The community attributes are updated by default.

If this command is configured for a neighbor group or address family group, all neighbors using the group inherit the configuration. Configuring the command specifically for a neighbor overrides inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to disable the router that sends community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/RSP0/CPU0:router(config)#router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp
```

The following example shows how to disable the delivery of community attributes to neighbor 172.20.1.1, preventing this feature from being inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)#router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# send-community-ebgp
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# send-community-ebgp inheritance-disable
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-extended-community-ebgp, on page 262 | Specifies that extended community attributes are sent to eBGP neighbors. |

send-extended-community-ebgp

To specify that extended community attributes should be sent to external Border Gateway Protocol (eBGP) neighbors, use the **send-extended-community-ebgp** command in an appropriate configuration mode. To disable sending extended community attributes to eBGP neighbors, use the **no** form of this command.

send-extended-community-ebgp [**inheritance-disable**]

no send-extended-community-ebgp [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|--|
| inheritance-disable | (Optional) Allows configurations inherited from a neighbor group or address family group to be overridden. |
|----------------------------|--|

Command Default

Extended community (EXTCOMM) attributes are NOT sent to eBGP peers (including PE-CE peers).

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 IPv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration
 VPNv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **send-extended-community-ebgp** command to control whether extended community attributes are sent to eBGP neighbors. This command cannot be used for iBGP neighbors as extended community attributes are always sent to iBGP neighbors.

When IOS XR BGP updates community attributes for eBGP VPN peers (VPNv4 or VPNv6), there is no need to configure the **send-extended-community-ebgp** command separately. The community attributes are updated by default.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure the router to send extended community attributes to neighbor 172.20.1.1 for IP Version 4 (IPv4) multicast routes:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp
```

The following example shows how to disable the delivery of extended community attributes to neighbor 172.20.1.1, preventing this feature from being automatically inherited from address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 140
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# send-extended-community-ebgp
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# send-extended-community-ebgp inheritance-disable
```

Related Commands

| Command | Description |
|--|---|
| af-group, on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| send-community-ebgp, on page 260 | Specifies that community attributes should be sent to an eBGP neighbor. |

session-group

To create a session group and enter session group configuration mode, use the **session-group** command in router configuration mode. To remove a session group and delete all configurations associated with it, use the **no** form of this command.

session-group *name*

no session-group *name*

Syntax Description

| | |
|-------------|----------------------------|
| name | Name of the session group. |
|-------------|----------------------------|

Command Default

No session groups are created.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **session-group** command to create a session group from which neighbors can inherit configuration that is address family-independent. That is, session groups cannot have address family-specific configuration. This command enters the session group configuration mode in which configuration for a session group is entered.

Many commands can be configured in both session group configuration mode and neighbor configuration mode.

Use of session groups saves time and reduces the router configuration size. Because the configuration of a session group can be inherited by any number of neighbors, use of the group can eliminate the need to copy long or complex configurations on each of a large number of neighbors. A neighbor can inherit all configuration from a session group simply by configuring the **use** command. Specific inherited session group configuration commands can be overridden for a specific neighbor by explicitly configuring the command for the specific neighbor.

The **no** form of this command causes all of the configuration for the session group to be removed. You cannot use the **no** form of this command if removing the group would leave one or more neighbors without a configured remote autonomous system number.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows a session group called group1 that is used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group.

```
RP/0/RSP0RP0/CPU0:router(config)# router bgp 1
RP/0/RSP0RP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0RP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0RP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0RP0/CPU0:router(config-bgp-nbr)# use session-group group1
```

The following example shows a session group called group1 used by two neighbors, 10.0.0.1 and 10.0.0.2. Because group1 is a session group, it contains only address family-independent configuration. And because group1 is used by neighbors 10.0.0.1 and 10.0.0.2, they inherit the configuration of the group. However, the **password password1** configuration from group1 is overridden for neighbor 10.0.0.2, using the **password-disable** command in the neighbor 10.0.0.2 configuration submode.

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 2
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# password password1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.0.0.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# password-disable
```

session-open-mode

To establish a Border Gateway Protocol (BGP) session with a specific TCP open mode, use the **session-open-mode** command in an appropriate configuration mode. To restore the default state, use the **no** form of this command.

session-open-mode {**active-only**| **both**| **passive-only**}

no session-open-mode [**active-only**| **both**| **passive-only**]

Syntax Description

| | |
|---------------------|---|
| active-only | Ensures that the BGP session can be established only when the request is initiated by the local end (active-open request) and all passive-open requests (from the other end) are rejected by the local BGP. |
| both | Allows BGP sessions to be established from both incoming or outgoing TCP connection requests, with one being rejected in the event of a request collision. |
| passive-only | Ensures that the local BGP does not initiate any TCP open requests and the session can be established only when the request comes from the remote end. |

Command Default

The default is **both**.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

BGP, by default, tries to initiate an active TCP connection whenever a new neighbor is configured. A remote neighbor may also initiate the TCP connection before the local BGP can initiate the connection. This initiation of a TCP connection by a remote neighbor is considered a passive-open request and it is accepted by the local BGP. This default behavior can be modified using the **session-open-mode** command.

**Note**

The BGP connection is not opened and, as a result the BGP session, is not established if both the peering neighbors use the same nondefault TCP session open mode—active-only or passive-only. If both ends are configured with active-only, each neighbor rejects the TCP open request from the other end. One neighbor must be configured as passive-only or both. Similarly, if both neighbors are configured with passive-only, neither neighbor initiates the TCP open request and the BGP session is not established. Again, one neighbor must be configured as active-only or both. There is one exception. A connection open request from a neighbor that is configured with the TCP session open mode to be passive-only is processed to detect whether there is a connection collision before the request is rejected. This exception enables the local BGP to reset the session if the remote neighbor goes down and it is not detected by the local router.

Use the **session-open-mode** command when it may be necessary to preconfigure a neighbor that does not exist. Ensure that BGP does not spend any time actively trying to set up a TCP session with the neighbor. A BGP session does not come up between two neighbors, both of which configure the same nondefault value (**active-only** or **passive-only** keyword) for this command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable a BGP session on router bgp 1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 45.67.89.01
RP/0/RSP0/CPU0:router(config-bgp-nbr)# session-open-mode active-only
```

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel}| ipv6 {unicast| multicast| all|
labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast [rd rd-address]]
vrf {vrf-name| all} [ipv4 {unicast| labeled-unicast}| ipv6 unicast]] [rd rd-address]] [ip-address [{mask|
/prefix-length} [longer-prefixes| unknown-attributes| bestpath-compare]]] [standby] [detail]
```

Syntax Description

| | |
|---|--|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd <i>rd-address</i> | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| <i>ip-address</i> | (Optional) Network address, entered to display a particular network in the BGP routing table. If the network address is omitted, then all networks in the BGP routing table are displayed. If the network mask and prefix length is omitted, then the software displays the longest matching prefix for the network address. |
| <i>mask</i> | (Optional) Network mask of the BGP route to match. |

| | |
|---------------------------|--|
| <i>/prefix-length</i> | (Optional) Prefix length of the BGP route to match. A slash (/) must precede the decimal value. |
| longer-prefixes | (Optional) Displays a route with the specified prefix length and more-specific routes if available. The longer-prefixes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| unknown-attributes | (Optional) Includes unknown, transitive attributes. The unknown-attributes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| bestpath-compare | (Optional) Displays route and best-path comparison information. The bestpath-compare keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.1.1 | The command output was modified to display from BGP Accept Own configuration. |
| Release 4.0.0 | The command output was modified to display BGP add-path information. |
| Release 4.3.0 | The command output was modified to include information from update wait-install configuration. |
| Release 5.1.1 | The command output was modified to display the status of permanent paths. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the sessions and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for an address family or a subaddress family, each matching routing table is examined in turn.

**Note**

Running the **show bgp** command immediately after configuring a large and complex route policy may result in timeout of the system database shown through an error message (SYSDB-SYSDB-6-TIMEOUT_EDM). It is recommended, that the show command be run, after the new route policy takes effect.

Use the **show bgp ip-address { mask | / prefix-length }** command to display detailed information for a specific route. If the mask and prefix length are omitted, the details of the longest matching prefix for the IP address are displayed.

Use the **show bgp** command to display all routes in the specified BGP routing table. Use the **show bgp ip-address { mask | / prefix-length } longer-prefixes** command to display those routes more specific than a particular prefix.

Use the **unknown-attributes** keyword to display details of any transitive attributes associated with a route that are not understood by the local system.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp** command in EXEC mode:

```
RP/0/RSP0/CPU0:router#show bgp
BGP router identifier 172.20.1.1, local AS number 1820
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 3
Dampening enabled
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop           Metric LocPrf Weight Path
* i10.3.0.0/16  172.20.22.1        0      100     0 1800 1239 ?
*>i            172.20.16.1        0      100     0 1800 1239 ?
* i10.6.0.0/16  172.20.22.1        0      100     0 1800 690 568 ?
*>i            172.20.16.1        0      100     0 1800 690 568 ?
```

```

* i10.7.0.0/16      172.20.22.1      0 100      0 1800 701 35 ?
*>i                172.20.16.1      0 100      0 1800 701 35 ?
*                  192.168.40.24    0 100      0 1878 704 701 35 ?
* i10.8.0.0/16     172.20.22.1      0 100      0 1800 690 560 ?
*>i                172.20.16.1      0 100      0 1800 690 560 ?
*                  192.168.40.24    0 100      0 1878 704 701 560 ?
* i10.13.0.0/16    172.20.22.1      0 100      0 1800 690 200 ?
*>i                172.20.16.1      0 100      0 1800 690 200 ?
*                  192.168.40.24    0 100      0 1878 704 701 200 ?
* i10.15.0.0/16    172.20.22.1      0 100      0 1800 174 ?
*>i                172.20.16.1      0 100      0 1800 174 ?
* i10.16.0.0/16    172.20.22.1      0 100      0 1800 701 i
*>i                172.20.16.1      0 100      0 1800 701 i
*                  192.168.40.24    0 100      0 1878 704 701 i

```

Processed 8 prefixes, 8 paths

This table describes the significant fields shown in the display.

Table 3: show bgp Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between BGP scans for the specified address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric. |

| Field | Description |
|--------|---|
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the path origin code. |

The following is sample output from the **show bgp** command with the network specified:

```
RP/0/RSP0/CPU0:router# show bgp 11.0.0.0/24
BGP router table entry for 11.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          2         2

Paths: (3 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    10.4.101.1
  Received by speaker 0
  Local
    0.0.0.0 from 0.0.0.0 (10.4.0.1)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, local, best
  Received by speaker 0
  2 3 4
    10.4.101.1 from 10.4.101.1 (10.4.101.1)
      Origin IGP, localpref 100, valid, external
  Received by speaker 0
  Local
    10.4.101.2 from 10.4.101.2 (10.4.101.2)
      Origin IGP, localpref 100, valid, internal
```

This table describes the significant fields shown in the display.

Table 4: show bgp prefix length Field Descriptions

| Field | Description |
|------------------------|--|
| BGP router table entry | Network that is being displayed. |
| Versions | List of the network versions in each BGP process. |
| Process | Name of the BGP process. |
| bRIB/RIB | Version of the network for sending to the RIB. You can compare this version with the bRIB/RIB version for the process (at the top of show bgp summary) to verify whether the network has been sent to the RIB. |

| Field | Description |
|---|---|
| SendTblVer | Version of the network for advertising to neighbors. This can be compared with the neighbor version to determine whether the network has been advertised to a particular neighbor. |
| Paths | List of paths for the network (that is, routes to reach the network). The number of paths and the index of the best path are given. |
| not advertised to any peer | Best path was received with a NO_ADVERTISE community and is not advertised to any neighbor. |
| not advertised to EBGp peer | Best path was received with a NO_EXPORT community and is not advertised to any eBGp neighbor. |
| not advertised outside local AS | Best path was received with a LOCAL_AS community and is not advertised to peers outside the local AS. |
| Advertisements of this net are suppressed by an aggregate | Network is a more-specific prefix of a configured aggregate and has been suppressed. It is not advertised to any neighbors unless they have an unsuppress-map configured. |
| Advertised to update-groups | List of update-groups to which the net has been advertised. Update-groups that have only one peer are not listed here. |
| Advertised to peers | List of neighbors to which the net has been advertised to. Neighbors that are in one of the update-groups listed above are not listed separately. Only neighbors that are in unique update-groups are listed. |
| Received by speaker 0 | BGP process where the path originated. This is always "speaker 0" for standalone mode. It will be the speaker-id when BGP is in distributed mode. |
| AS Path | Autonomous system (AS) path that was received for the path. If the AS path is empty, then "Local" is displayed. This is the case for paths that are locally generated on this router or on a neighboring router within the same AS. |
| aggregated by | If the path is an aggregate, the router-id of the router that performed the aggregation. |
| suppressed due to dampening | Path has been suppressed due to the configured path dampening. |

| Field | Description |
|---|--|
| history entry | Path is withdrawn, but a copy is kept to store the dampening information. |
| Received from a RR-client | Path was received from a route reflector client. |
| received-only | If soft reconfiguration inbound is configured, the path was received but dropped by inbound policy, or was accepted and modified. In either event, the received-only value is a copy of the original, unmodified path. |
| received & used | If soft reconfiguration inbound is configured, the path was received and accepted by inbound policy, but not modified. |
| stale | Neighbor from which the path was received is down, and the path is kept and marked as stale to support graceful restart. |
| <nexthop> from <neighbor> (<router-id>) | Next hop for the path. If the next hop is known by a mechanism outside BGP (for example, for redistributed paths), then 0.0.0.0 is displayed. After the next hop, the neighbor from whom the path was received is displayed, along with the neighbor's router-id. If the path was locally generated (for example, an aggregate or redistributed path), then 0.0.0.0 is displayed for the neighbor address. |
| Origin | IGP: the path originated from an IGP. EGP: the path originated from an EGP. incomplete: the origin of the path is unknown. |
| metric | MED value of the path. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| weight | Locally assigned weight (if not 0) of the path. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| valid | Path is valid and can be considered in the best-path calculation. |
| redistributed | Path is redistributed through a redistribute command. |

| Field | Description |
|-----------------------|---|
| aggregated | Path is a locally generated aggregate created due to an aggregate-address command. |
| local | Path is a local network source due to a network command. |
| internal | Path was received from an iBGP neighbor. |
| external | Path was received from an eBGP neighbor. |
| atomic-aggregate | Path was received with the atomic-aggregate flag set. Some path information has been removed through aggregation. |
| best | Path is the best path for the network and is used for routing and advertised to peers. |
| multipath | Path is a multipath and is installed into the RIB along with the best path. |
| Community | List of communities attached to the path. |
| Extended community | List of extended communities attached to the path. |
| Originator | Originator of the path within the AS Cluster list if the path is reflected. |
| AS Cluster list | List of RR clusters the path has passed through if the path is reflected |
| Dampinfo | Penalty and reuse information if the path is dampened. |
| penalty | Current penalty for the path. |
| flapped | Number of times the path has flapped and the time since the first flap. |
| reuse in | Time until the path is re-used (undampened). |
| half life | Configured half-life for the path. |
| suppress value | Penalty at which the path is suppressed. |
| reuse value | Penalty at which the path is re-used. |
| Maximum suppress time | Maximum length of time for which the path can be suppressed. |

The following is sample output from the **show bgp** command with the *ip-address/prefix-length detail* options:

```
RP/0/RSP0/CPU0:router# show bgp 51.0.0.0/24 detail
Sat Mar 14 00:37:14.109 PST PDT
BGP routing table entry for 51.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          3         3
  Flags: 0x3e1000, label_retention: not enabled
Last Modified: Mar 13 19:32:17.976 for 05:04:56
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.3 0.4 0.7 0.8
  Advertised to peers (in unique update groups):
    201.48.20.1
  Path #1: Received by speaker 0
  Flags: 0x1000003
  200 201
    213.0.0.6 from 213.0.0.6 (200.200.3.1)
    Origin IGP, localpref 100, valid, external, best
```

The following is sample output from the **show bgp** command with the additional paths received from:

```
BGP routing table entry for 51.0.1.0/24, Route Distinguisher: 2:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          63        63
  Flags: 0x040630f2
Last Modified: Nov 11 12:44:05.811 for 00:00:16
Paths: (3 available, best #2)
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  Path #1: Received by speaker 0
  Flags: 0x3
  Not advertised to any peer
  111 111 111 111 111 111 111 111
    10.51.0.10 from 10.51.0.10 (11.11.11.11)
    Origin IGP, metric 0, localpref 100, valid, external
    Received Path ID 0, Local Path ID 0, version 0
    Extended community: RT:55:1
  Path #2: Received by speaker 0
  Flags: 0x5060007
  Advertised to CE peers (in unique update groups):
    10.51.0.10
  561 562 563 564 565
    13.0.6.50 from 13.0.6.50 (13.0.6.50)
    Received Label 16
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
    Received Path ID 0, Local Path ID 1, version 63
    Extended community: RT:55:1
  Path #3: Received by speaker 0
```

```

Flags: 0x4060007
Not advertised to any peer
591 592 593 594 595
  13.0.9.50 from 13.0.9.50 (13.0.9.50)
    Received Label 16
    Origin IGP, localpref 100, valid, internal, backup, add-path, import-candidate,
imported
    Received Path ID 0, Local Path ID 4, version 63
    Extended community: RT:22:232 RT:55:1

```

This is sample output to explain 'import suspect' state and 'import-suspect' field in **show bgp** command output:

```

RP/0/RSP0/CPU0:router#show bgp vpnv4 unicast rd 11:111 100.16.11.0/24
BGP routing table entry for 100.16.11.0/24, Route Distinguisher: 11:111
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          1834195   1834195
Paths: (2 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Path #1: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 55.55.55.55 (16.16.16.16)
      Received Label 19602
      Origin incomplete, localpref 100, valid, internal, best, import-candidate, not-in-vrf,
import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 55.55.55.55
  Path #2: Received by speaker 0
  11
    1:16.16.16.16 (metric 30) from 88.88.88.88 (16.16.16.16)
      Received Label 19602
      Origin incomplete, localpref 100, valid, internal, not-in-vrf, import suspect
      Extended community: RT:11:11
      Originator: 16.16.16.16, Cluster list: 88.88.88.88

```

The **show bgp** command output displays 'import suspect' when potential import oscillation has been detected for the prefix. Import of such a prefix is not affected. However, import of the prefix can be dampened in future if the oscillation continues. If the oscillation stops during the next import run, the prefix will no longer be marked 'import suspect'.

This is sample output of **show bgp {ipv4 | vpnv4} unicast summary** when the **update wait-install** command was configured for an address family. The output displays the "RIBAckVer" field.

```

RP/0/RSP0/CPU0:router#show bgp summary

BGP router identifier 10.1.1.2, local AS number 100
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000   RD version: 5
BGP main routing table version 5
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process  RcvTblVer   RIBVer/RIBAckVer   LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker          5             5/5                5          5           5            5

Neighbor      Spk   AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
10.1.1.1       0    500    0        0         0     0    0 00:00:00  Idle

```

This is sample output from **show bgp vpnv4 unicast rd prefix/length** command that displays Accept Own prefix information:

```

RP/0/RSP0/CPU0:router#show bgp vpnv4 unicast rd 10.10.10.10:1 110.1.1.1/32 detail
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 10.10.10.10:1
Versions:

```

```

Process          bRIB/RIB  SendTblVer
Speaker          1412487   1412487
  Local Label: 137742 (no rewrite);
  Flags: 0x04043001+0x00000000;
Last Modified: Jul 19 14:42:43.690 for 00:56:34
Paths: (2 available, best #1)
  Advertised to peers (in unique update groups):
    45.1.1.1
  Path #1: Received by speaker 0
  Flags: 0xd040003, import: 0x1f
  Advertised to peers (in unique update groups):
    45.1.1.1
101
  10.5.1.2 from 10.5.1.2 (10.5.1.2)
    Origin incomplete, localpref 100, valid, external, best, group-best, import-candidate

    Received Path ID 0, Local Path ID 1, version 1412487
    Extended community: RT:100:1
  Path #2: Received by speaker 0
  Flags: 0x324020005, import: 0x01
  Not advertised to any peer
101
  15.1.1.1 from 55.1.1.1 (15.1.1.1)
    Received Label 137742
    Origin incomplete, localpref 100, valid, internal, import-candidate, not-in-vrf,
accept-own-self
    Received Path ID 0, Local Path ID 0, version 0
    Community: accept-own
    Extended community: RT:100:1 RT:1000:1
    Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1

```

This is sample output from **show bgp vrf vrf-name ipv4unicast prefix/length** command that displays Accept Own prefix information on a customer (originating) VRF:

```

RP/0/RSP0/CPU0:router#show bgp vrf customer1 ipv4 uni 110.1.1.1/32
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 10.10.10.10:1
Versions:
Process          bRIB/RIB  SendTblVer
Speaker          1412487   1412487
  Local Label: 137742
Last Modified: Jul 19 14:42:43.690 for 01:01:22
Paths: (2 available, best #1)
  Advertised to PE peers (in unique update groups):
    45.1.1.1
  Path #1: Received by speaker 0
  Advertised to PE peers (in unique update groups):
    45.1.1.1
101
  10.5.1.2 from 10.5.1.2 (10.5.1.2)
    Origin incomplete, localpref 100, valid, external, best, group-best, import-candidate

    Received Path ID 0, Local Path ID 1, version 1412487
    Extended community: RT:100:1
  Path #2: Received by speaker 0
  Not advertised to any peer
101
  15.1.1.1 from 55.1.1.1 (15.1.1.1)
    Received Label 137742
    Origin incomplete, localpref 100, valid, internal, import-candidate, not-in-vrf,
accept-own-self
    Received Path ID 0, Local Path ID 0, version 0
    Community: accept-own
    Extended community: RT:100:1 RT:1000:1
    Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1

```

This is sample output from **show bgp vrf vrf-name ipv4unicast prefix/length** command that displays Accept Own prefix information on a service VRF:

```

RP/0/RSP0/CPU0:router#show bgp vrf servicel ipv4 uni 110.1.1.1/32
BGP routing table entry for 110.1.1.1/32, Route Distinguisher: 11.11.11.11:1

```

show bgp

```

Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          1412497   1412497
Last Modified: Jul 19 14:43:08.690 for 01:39:22
Paths: (1 available, best #1)
  Advertised to CE peers (in unique update groups):
    10.8.1.2
  Path #1: Received by speaker 0
  Advertised to CE peers (in unique update groups):
    10.8.1.2
  101
  10.5.1.2 from 55.1.1.1 (15.1.1.1)
  Origin incomplete, localpref 100, valid, internal, best, group-best, import-candidate,
  imported, accept-own
  Received Path ID 0, Local Path ID 1, version 1412497
  Community: accept-own
  Extended community: RT:100:1 RT:1000:1
  Originator: 15.1.1.1, Cluster list: 55.1.1.1, 75.1.1.1, 45.1.1.1

```

This table describes the significant fields shown in the display:

| Field | Description |
|----------------------|--|
| accept-own-self | The Accept Own path in the customer VRF contains the "accept-own-self" keyword/flag. |
| accept-own | The Accept Own path contains the "accept-own" keyword/flag. |
| Community:accept-own | List of communities attached to the path: accept-own. |
| Extended community | List of extended communities attached to the path. |
| Cluster list | Router ID or cluster ID of all route reflectors through which the route has passed. |

The output of **show bgp {vpn4 | vpn6} unicast rd** command may display the optional BGP attribute `not-in-vrf`. If a path in a VPNvX net is marked as `not-in-vrf`, it may be due to any of the following conditions:

- The RD of the VPNvX net is not the same as any of the RDs configured for VRFs on the router.
- The RD of the VPNvX net is the same as the RD configured for a specific VRF on the router, but the path is not imported to the specified VRF. For example, the route-targets attached to the path do not match any of the **import route-target** [*as-number:nn* | *ip-address:nn*] configured for VRF, *vrf_1*.

If the `not-in-vrf` net is set, it indicates that the path does not belong to the VRF.

This is sample output from the **show bgp ipv4 unicast** command showing the status of the permanent network:

```

RP/0/RSP0/CPU0:router# show bgp ipv4 unicast 1.0.0.0/24
BGP routing table entry for 1.0.0.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          90113     90113
Last Modified: Sep  6 04:46:03.650 for 00:14:19
Permanent Network
Paths: (2 available, best #2)
  Advertised to peers (in unique update groups):
    2.2.2.2
  Path #1: Received by speaker 0

```

```

Advertised to peers (in unique update groups):
 3.3.3.3
Local
 0.0.0.0 from 0.0.0.0 (1.1.1.1)
  Origin incomplete, metric 0, localpref 100, local, permanent-path
  Received Path ID 0, Local Path ID 4, version 90113
  Origin-AS validity: not-found
Path #2: Received by speaker 0
  Advertised to peers (in unique update groups):
  2.2.2.2
  7813 7814
  11.11.22.22 from 11.11.22.22 (192.1.1.1)
  Origin EGP, localpref 100, valid, external, best, group-best, import-candidate
  Received Path ID 0, Local Path ID 1, version 4
  Origin-AS validity: not-found

```

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 28 | Creates an aggregate entry in a BGP routing table. |
| bgp default local-preference , on page 78 | Changes the default local preference value. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |
| show bgp cidr-only , on page 309 | Displays routes with nonnatural netmasks. |
| show bgp community , on page 313 | Displays routes belonging to the specified communities. |
| show bgp inconsistent-as , on page 332 | Displays networks with inconsistent origin autonomous system. |
| show bgp regexp , on page 426 | Displays routes matching an AS path regular expression. |
| show bgp route-policy , on page 431 | Displays networks that match a route policy. |
| show bgp summary , on page 442 | Displays the status of all BGP connections. |
| show bgp truncated-communities , on page 451 | Displays networks with community lists truncated by policy. |

show bgp update out

To display address-family level update generation information, use the **show bgp update out** command in EXEC mode.

show bgp [*vrf vrf-name*] [*afi safi*] **update out** [**brief**] **detail**

Syntax Description

| | |
|----------------------------|--|
| vrf <i>vrf-name</i> | (Optional) Displays non-default VRF. |
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

Examples

This example displays sample output from the **show bgp update out** command:

```
RP/0/RSP0/CPU0:router#show bgp update out
Address-family "IPv4 Unicast"
  Update generation status: Normal
  Update OutQ:                0 bytes (0 messages)
  AF update limit: 268435456 bytes (configured 268435456 bytes)
```

```
EBGP Sub-group update limit: 33554432 bytes (configured 33554432 bytes)
IBGP Sub-group update limit: 33554432 bytes (configured 33554432 bytes)

Main routing table version: 2
RIB version: 2
Minimum neighbor version: 2
AF Flags: 0x00000000
Update-groups: 1
Sub-groups: 1 (0 throttled)
Refresh sub-groups: 0 (0 throttled)
Filter-groups: 1
Neighbors: 3

History:
  Update OutQ Hi: 300 bytes (1 messages)
  Update OutQ Cumulative: 600 bytes (2 messages)
  Update OutQ Discarded: 0 bytes (0 messages)
  Update OutQ Cleared: 0 bytes (0 messages)
  Last discarded from OutQ: --- (never)
  Last cleared from OutQ: --- (never)
  Update generation throttled 0 times, last event --- (never)
  Update generation recovered 0 times, last event --- (never)
  Update generation mem alloc failed 0 times, last event --- (never)

VRF "default", Address-family "IPv4 Unicast"
  RD flags: 0x00000001
  RD Version: 2
  Table flags: 0x00000021
  RIB version: 2
  Update-groups: 1
  Sub-groups: 1 (0 throttled)
  Refresh sub-groups: 0 (0 throttled)
  Filter-groups: 1
  Neighbors: 3

RP/0/RSP0/CPU0:PE51_ASR-9010#
RP/0/RSP0/CPU0:PE51_ASR-9010#
RP/0/RSP0/CPU0:PE51_ASR-9010#show bgp update out filter-group
Thu Sep 13 01:43:48.183 DST
```

show bgp update in error process

To display process level update inbound error-handling information, use the **show bgp update in error process** command in EXEC mode.

show bgp update in error process [**brief**] **detail**

| Syntax Description | |
|--------------------|--|
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default None

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples This example displays sample output from the **show bgp update in error process** command:

```
RP/0/RSP0/CPU0:router#show bgp update in error process
```

```
Basic Update error-handling:
  EBGP: [Enabled]
  IBGP: [Enabled]
Extended Update error-handling:
  EBGP: [Disabled]
  IBGP: [Disabled]
```

```
Malformed Update messages: 0
Neighbors that received malformed Update messages: 0
Last malformed Update received: --- (never)
```

show bgp update out filter-group

To display update generation information at filter-group level, **show bgp update out filter-group** command in EXEC mode.

show bgp [*vrf vrf-name*] [*afi safi*] **update out filter-group** [*fg-process-id*] [**brief**] **detail**

Syntax Description

| | |
|----------------------------|--|
| vrf <i>vrf-name</i> | Specifies the non-default VRF. |
| <i>afi safi</i> | Specifies the address family and subsequent address family identifiers. |
| <i>fg-process-id</i> | Specifies the filter-group process ID in <x.y> format. Range is <0-15>.<0-4294967295>. |
| brief | (Optional) Displays brief information on filter-group level update generation |
| detail | (Optional) Displays detailed information on filter-group level update generation. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|-----------------------------|
| Release 4.2.0 | This command was introduced |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

This example displays sample output from **show bgp update out filter-group** command:

show bgp update out process

To display process level update generation information, use the **show bgp update out process** command in EXEC mode.

show bgp update out process [**brief**] **detail**

| Syntax Description | |
|--------------------|--|
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |

Command Default None

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-----------|
| | bgp | read |

Examples This example displays sample output from the show bgp update out process brief command:

```
RP/0/RSP0/CPU0:router#show bgp update out process
Wed Sep 12 08:26:04.308 DST

Update generation status: Normal
Update OutQ:           0 bytes (0 messages)
Update limit:  536870912 bytes (configured 536870912 bytes)

Update generation logging:  [Disabled]

  Address-family Status   Limit      OutQ      UG   SG(Thr)   SG-R(Thr) Nbrs
  -----
  IPv4 Unicast   Normal    268435456  0         1     1(0)     0(0)      3
  L2VPN VPLS    Normal    268435456  0         1     1(0)     0(0)      3
```

```

History:
  Update OutQ Hi:                300 bytes (1 messages)
  Update OutQ Cumulative:        1200 bytes (4 messages)
  Update OutQ Discarded:         0 bytes (0 messages)
  Update OutQ Cleared:           0 bytes (0 messages)
  Last discarded from OutQ: --- (never)
  Last cleared from OutQ: --- (never)
  Update generation throttled 0 times, last event --- (never)
  Update generation recovered 0 times, last event --- (never)
  Update generation mem alloc failed 0 times, last event --- (never)

```

This table describes the significant fields shown on the display:

Table 5: show bgp update out process Field Descriptions

| Field | Description |
|--------------------------|-------------|
| Update generation status | |
| Update OutQ | |

show bgp update out sub-group

To display sub-group update generation information, use the **show bgp update out sub-group** command in EXEC mode.

show bgp [*vrf vrf-name*] [*afi safi*] **update out** [*update-group ug-index*] **sub-group** [*sg-index*] [**brief**] **detail**

Syntax Description

| | |
|----------------------------|--|
| vrf <i>vrf-name</i> | (Optional) Displays non-default VRF. |
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |
| <i>ug-index</i> | (Optional) Displays the update-group process ID in <x.y> format. |
| <i>sg-index</i> | (Optional) displays the sub-group process ID in <x.y> format. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

Examples

This example displays sample output from the **show bgp update out sub-group** command:

```
RP/0/RSP0/CPU0:router#show bgp update out sub-group
VRF "default", Address-family "IPv4 Unicast"
  Main routing table version: 2
  RIB version: 2

  SG          UG      Status   Limit    OutQ      SG-R Nbrs  Version   ()
  0.2         0.2      Normal   33554432  0         0   3   2         ()
RP/0/RSP0/CPU0:PE51_ASR-9010#
```

This table describes the significant fields shown in the display:

show bgp update out update-group

To display update-group update generation information, use the **show bgp update out update-group** command in EXEC mode.

show bgp [*vrf vrf-name*] [*afi safi*] **update out update-group** [*ug-index*] [**brief**] **detail**

Syntax Description

| | |
|----------------------------|--|
| vrf <i>vrf-name</i> | (Optional) Displays non-default VRF. |
| <i>afi</i> | (Optional) Displays address-family identifier. |
| <i>safi</i> | (Optional) Displays subsequent address family identifier. |
| brief | (Optional) Displays brief information on process level update generation. |
| detail | (Optional) Displays detailed information on process level update generation. |
| <i>ug-index</i> | (Optional) Displays the update-group process ID in <x.y> format. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

Examples

This example shows the significant fields on display form the **show bgp update out update-group** command:

```
RP/0/RSP0/CPU0:router#show bgp update out sub-group
```

```

VRF "default", Address-family "IPv4 Unicast"
  Main routing table version: 2
  RIB version: 2

  SG          UG      Status   Limit      OutQ      SG-R Nbrs Version  ()
  0.2         0.2     Normal  33554432   0          0   3   2          ()
RP/0/RSP0/CPU0:PE51_ASR-9010#show bgp update ou update-group
Wed Sep 12 08:37:24.756 DST

```

```

VRF "default", Address-family "IPv4 Unicast"

  UG      OutQ      SG(Thr)   SG-R(Thr)  FG   Nbrs
  0.2     0           1(0)      0(0)       1    3

```

This table describes the significant fields on the display:

show bgp vrf update in error

To display VRF level update inbound error-handling information, use the **show bgp vrf update in error** command in EXEC mode.

show bgp [*vrf vrf-name*] **update in error** [**brief**] **detail**

| | |
|----------------------------|---|
| vrf <i>vrf-name</i> | (Optional) Displays non-default VRF. |
| brief | (Optional) Displays brief information. |
| detail | (Optional) Displays detailed information. |

Command Default None

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read |

Examples This example displays sample output from **show bgp vrf vrf1 update in error** command:

```
RP/0/RSP0/CPU0:router#show bgp update in error

VRF "default"
  Malformed Update messages: 0
  Neighbors that received malformed Update messages: 0
  Last malformed update received: --- (never)
```

show bgp advertised

To display advertisements for neighbors or a single neighbor, use the **show bgp advertised** command in EXEC mode.

```
show bgp [ipv4 { all | labeled-unicast | mdt | multicast | tunnel | unicast } ] advertised [neighbor ip-address]
[standby] [summary]
```

```
show bgp [ ipv6 { all | labeled-unicast | multicast | unicast } ] advertised [neighbor ip-address] [standby]
[summary]
```

```
show bgp [ all { all | labeled-unicast | multicast | tunnel | unicast } ] advertised [neighbor ip-address]
[standby] [summary]
```

```
show bgp [ vpnv4 unicast [rd rd-address]] advertised [neighbor ip-address] [standby] [summary]
```

```
show bgp [ vpnv6 unicast [ rd rd-address]] advertised [neighbor ip-address] [standby] [summary]
```

```
show bgp [ vrf {vrf-name| all} [ ipv4 | {labeled-unicast | unicast} | ipv6 unicast]] advertised [neighbor
ip-address] [standby] [summary]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |

| | |
|---------------------|---|
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| neighbor | (Optional) Previews advertisements for a single neighbor. If the neighbor keyword is omitted, then the advertisements for all neighbors are displayed. |
| <i>ip-address</i> | (Optional) IP address of the neighbor. |
| summary | (Optional) Displays a summary of advertisements. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that is configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp advertised** command to display the routes that have been advertised to peers or a specific peer. To preview advertisements that would be sent to a peer under a particular policy, even if the corresponding update messages have not been generated yet, use the **show bgp policy** command.



Note When you issue the **show bgp advertised** command, a route is not displayed in the output unless an advertisement for that route has already been sent (and not withdrawn). If an advertisement for the route has not yet been sent, the route is not displayed.

Use the **summary** keyword to display a summary of the advertised routes. If you do not specify the **summary** keyword, the software displays detailed information about the advertised routes.



Note The **show bgp advertised** command does not display the application of any outbound policy in the route details it displays. Consequently, this command provides only an indication of whether a particular route has been advertised, rather than details of which attributes were advertised. Use the **show bgp policy sent-advertisements** command to display the attributes that are advertised.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp advertised** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp advertised neighbor 10.0.101.4 summary
Network      Next Hop      From           AS Path
1.1.1.0/24   10.0.101.1   10.0.101.1    2 3 222 333 444 555 i
1.1.2.0/24   10.0.101.1   10.0.101.1    3 4 5 6 7 i
1.1.3.0/24   10.0.101.1   10.0.101.1    77 88 33 44 55 99 99 99 i
1.1.4.0/24   10.0.101.1   10.0.101.1    2 5 6 7 8 i
1.1.7.0/24   10.0.101.1   10.0.101.1    3 5 i
1.1.8.0/24   10.0.101.1   10.0.101.1    77 88 99 99 99 i
```

This table describes the significant fields shown in the display.

Table 6: show bgp advertised neighbor summary Field Descriptions

| Field | Description |
|----------|--|
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| From | IP address of the peer that advertised this route. |
| AS Path | AS path of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |

| Field | Description |
|-----------------|---|
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the peer to which this entry was advertised. This field is used in the output when displaying a summary of the advertisements to all neighbors. |

The following is sample output from the **show bgp advertised** command for detailed advertisement information:

```
RP/0/RSP0/CPU0:router# show bgp advertised neighbor 172.72.77.1
172.16.0.0/24 is advertised to 172.72.77.1
  Path info:
    neighbor: Local          neighbor router id: 172.74.84.1
    valid redistributed best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    MET ORG AS
    origin: incomplete metric: 0
    aspath:
10.52.0.0/16 is advertised to 172.72.77.1
  Path info:
    neighbor: Local Aggregate neighbor router id: 172.74.84.1
    valid aggregated best
  Attributes after inbound policy was applied:
  next hop: 0.0.0.0
    ORG AGG ATOM
    origin: IGP aggregator: 172.74.84.1 (1)
    aspath:
```

This table describes the significant fields shown in the display.

Table 7: show bgp advertised neighbor Field Descriptions

| Field | Description |
|--------------------|--|
| is advertised to | IP address of the peer to which this route has been advertised. If the route has been advertised to multiple peers, the information is shown separately for each peer. |
| neighbor | IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system. |
| neighbor router id | BGP identifier for the peer, or the local system if the route originated on the local system. |

| Field | Description |
|-------------------------------------|--|
| Not advertised to any peer | Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers. |
| Not advertised to any EBGP peer | Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those external peers are part of the same confederation as the local router. |
| Not advertised outside the local AS | Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. |
| (Received from a RR-client) | Path was received from a route reflector client. |
| (received-only) | This path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked "received-only" indicates that either the path was dropped by inbound policy, or the path information was modified by inbound policy and a separate copy of the modified path is used for routing. |
| (received & used) | Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked "received and used," implies the path information was not modified by inbound policy. |
| valid | Path is valid. |
| redistributed | Path is locally sourced through redistribution. |
| aggregated | Path is locally sourced through aggregation. |
| local | Path is locally sourced through the network command. |
| confed | Path was received from a confederation peer. |
| best | Path is selected as best. |
| multipath | Path is one of multiple paths selected for load-sharing purposes. |

| Field | Description |
|---|--|
| dampinfo | Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the router noticed the first flap. Reuse in—Time (hours:minutes:seconds) after which the path is made available. This field is displayed only if the path is currently suppressed. |
| Attributes after inbound policy was applied | Displays attributes associated with the received route, after any inbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present. |
| origin | Origin of the path: IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command. EGP—Path originated from an Exterior Gateway Protocol. incomplete—Origin of the path is not clear. For example, a route that is redistributed into BGP from an IGP. |
| neighbor as | First autonomous system (AS) number in the AS path. |
| aggregator | Indicates that the path was received with the aggregator attribute. The autonomous system number and router-id of the system that performed the aggregation are shown. |

| Field | Description |
|--------------------|--|
| metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system |
| aspath | AS path associated with the route. |
| community | Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities: Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers. no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation with the local router. |
| Extended community | Extended community attributes associated with the path. For known extended community types, the following codes may be displayed: RT—Route target community SoO—Site of Origin community LB—Link Bandwidth community |
| Originator | Router ID of the originating router when route reflection is used. |
| Cluster lists | Router ID or cluster ID of all route reflectors through which the route has passed. |

Related Commands

| Command | Description |
|-----------------|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |

| Command | Description |
|---|---|
| set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |
| route-policy (BGP), on page 249 | Applies a route policy to incoming and outgoing routes. |
| rd, on page 227 | Filters routes using a prefix list. |
| show bgp policy, on page 393 | Displays information about BGP advertisements under a proposed policy. |
| show bgp policy, on page 393 sent-advertisements | Previews advertisements to peers, including details of advertised attributes. |

show bgp af-group

To display information about Border Gateway Protocol (BGP) configuration for address family groups, use the **show bgp af-group** command in EXEC mode.

```
show bgp af-group group-name {configuration [defaults] [nvgen]| inheritance| users}
```

Syntax Description

| | |
|----------------------|--|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the af-group, including any settings that have been inherited from af-groups used by this af-group. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |
| nvgen | (Optional) Displays output in the format of show running-config output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups from which this af-group inherits configuration settings. |
| users | Displays the neighbors, neighbor groups, and af-groups that inherit configuration from this af-group. |

Command Default

No default behavior or value

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp af-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of an af-group, taking into account any configuration that may be inherited from other af-groups through the **use af-group** command. The source of each command is shown.

If the **defaults** keyword is specified, all configuration for the af-group, including default values, is shown. Default configuration is identified in the show output. Use the **nvgen** keyword to display configuration

formatted in the style of the **show running-config** command. This output is suitable for cutting and pasting into configuration sessions.

Use the **show bgp af-group** command with the *group-name* **inheritance** argument and keyword to display the address family groups from which the specified af-group inherits configuration.

Use the **show bgp af-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and af-groups that inherit configuration from the specified af-group.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following af-group configuration is used in the examples:

```
af-group group3 address-family ipv4 unicast
remove-private-AS
soft-reconfiguration inbound
!
af-group group1 address-family ipv4 unicast
use af-group group2
maximum-prefix 2500 75 warning-only
default-originate
soft-reconfiguration inbound disable
!
af-group group2 address-family ipv4 unicast
use af-group group3
send-community-ebgp
send-extended-community-ebgp
capability orf prefix both
```

The following is sample output from the **show bgp af-group** command with the **configuration** keyword in EXEC mode. The source of each command is shown in the right column. For example, **default-originate** is configured directly on **af-group group1**, and the **remove-private-AS** command is inherited from af-group group2, which in turn inherits it from af-group group3.

```
RP/0/RSP0/CPU0:router# show bgp af-group group1 configuration

af-group group1 address-family ipv4 unicast
capability orf prefix both           [a:group2]
default-originate                    []
maximum-prefix 2500 75 warning-only  []
remove-private-AS                    [a:group2 a:group3]
send-community                        [a:group2]
send-extended-community               [a:group2]
```

The following is sample output from the **show bgp af-group** command with the **users** keyword:

```
RP/0/RSP0/CPU0:router# show bgp af-group group2 users

IPv4 Unicast: a:group1
```

The following is sample output from the **show bgp af-group** command with the **inheritance** keyword. This example shows that the specified af-group group1 directly uses the group2 af-group, which in turn uses the group3 af-group:

```
RP/0/RSP0RPO/CPU0:router# show bgp af-group group1 inheritance
```

```
IPv4 Unicast: a:group2 a:group3
```

Table 8: [show bgp af-group Field Descriptions, on page 303](#) describes the significant fields shown in the display.

This table describes the significant fields shown in the display.

Table 8: show bgp af-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified address family group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the configuration is disabled by default. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|--|
| af-group, on page 26 | Configures a BGP address family group. |
| show bgp neighbors, on page 353 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp neighbor-group, on page 349 | Displays information about configuration for neighbor groups. |
| use, on page 514 af-group | Configures an af-group to inherit the configuration of a specified af-group. |

show bgp attribute-key

To display all existing attribute keys, use the **show bgp attribute-key** command in EXEC mode.

```
show bgp {ipv4| ipv6| all| vpnv4 unicast| vrf} attribute-key
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| all | (Optional) For subaddress family, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp attribute-key** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp all all attribute-key

Address Family: IPv4 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 109
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop        AttrKey
*> 1.1.0.0/16      0.0.0.0         0x00000002
*> 10.0.0.0/16     0.0.0.0         0x00000002
*> 12.21.0.0/16    0.0.0.0         0x00000002
*> 194.3.192.1/32  10.0.101.1      0x00000009
*> 194.3.192.2/32  10.0.101.1      0x00000009
*> 194.3.192.3/32  10.0.101.1      0x00000009
*> 194.3.192.4/32  10.0.101.1      0x00000009
*> 194.3.192.5/32  10.0.101.1      0x00000009

Processed 8 prefixes, 8 paths

Address Family: IPv4 Multicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 15
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop        AttrKey
*> 194.3.193.2/32  10.0.101.1      0x00000009
*> 194.3.193.3/32  10.0.101.1      0x00000009
```

```

Processed 2 prefixes, 2 paths

Address Family: IPv6 Unicast
=====

BGP router identifier 10.0.0.1, local AS number 1
BGP generic scan interval 60 secs
BGP main routing table version 19
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
               Network      Next Hop      AttrKey
*> 2222::1111/128      2222::2      0x00000009
*> 2222::1112/128      2222::2      0x00000009

Processed 2 prefixes, 2 paths

```

This table describes the significant fields shown in the display.

Table 9: show bgp attribute-key Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans. |

| Field | Description |
|---|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| AttrKey | Key associated with the route attribute. |
| Processed <i>n</i> prefixes, <i>n</i> paths | Number of prefixes and number of paths processed for the table. |

Related Commands

| Command | Description |
|-------------------------|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp cidr-only

To display routes with nonnatural network masks, also known as classless interdomain routing (CIDR) routes, use the **show bgp cidr-only** command in EXEC mode.

```
show bgp [ipv4|vrf] cidr-only [standby]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies the IP Version 4 address family. |
| unicast | (Optional) Specifies the unicast address family. |
| multicast | (Optional) Specifies the multicast address family. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress family, specifies all subaddress families. |
| tunnel | (Optional) Specifies the tunnel address family. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRE, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRE, specifies IPv4 unicast or labeled-unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used. This command is applicable only for IPv4 prefixes. If the default address family is not IPv4, then the **ipv4** keyword must be used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for subaddress family, all subaddress family routing tables are examined.

The **show bgp cidr-only** command applies only for IPv4 prefixes. If the **ipv4** keyword is not specified and the default address family is not IPv4, the command is not available.

Use the **show bgp cidr-only** command to display CIDR routes. Routes that have their correct class (class A, B, or C) prefix length are not displayed.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp cidr-only** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp cidr-only

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 2589
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop        Metric      LocPrf      Weight      Path
*> 192.0.0.0/8  192.168.72.24    0           1878        ?
*> 192.168.0.0/16 192.168.72.30    0           108         ?
```

This table describes the significant fields shown in the display.

Table 10: show bgp cidr-only Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |

| Field | Description |
|-------------------|--|
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Entry originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |

| Field | Description |
|--------|---|
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| route-policy (BGP), on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 268 | Displays BGP routes. |

show bgp community

To display routes that have the specified Border Gateway Protocol (BGP) communities, use the **show bgp community** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast | all| tunnel}] community community-list [exact-match]
```

```
show bgp [ipv6 {unicast}] community community-list [exact-match]
```

```
show bgp [all {unicast| multicast| labeled-unicast | all| tunnel}] community community-list [exact-match]
```

```
show bgp [vpnv4 unicast [rd rd-address]] community community-list [exact-match]
```

```
show bgp [vrf {vrf-name| all} [ipv4 {unicast| labeled-unicast}| ipv6 unicast]] community community-list [exact-match]
```

Syntax Description

| | |
|---|--|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd <i>rd-address</i> | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| community | Specifies that only routes with communities specified by <i>community-list</i> is displayed. |

| | |
|-----------------------|---|
| <i>community-list</i> | <p>Between one and seven communities. Each community can be a number in the range from 1 to 4294967295, a community specified in AA:NN format, or one of the following well-known communities:</p> <p>local-AS—Well-known community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Well-known community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Well-known community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> <p>internet—Well-known community whose value is not defined in BGP RFC. IOS XR BGP uses a value of 0 for the internet community. Routes with this community are advertised to all peers without any restrictions.</p> <p>For the AA:NN format:</p> <p>AA—Range is 0 to 65535.</p> <p>NN—Range is 1 to 4294967295.</p> <p>Up to seven community numbers can be specified.</p> |
| exact-match | (Optional) Displays those routes that have communities exactly matching the specified communities. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or the subaddress family, each matching routing table is examined in turn.

If more than seven communities are required, it is necessary to configure a route policy and use the [show bgp route-policy, on page 431](#) command.

Use the **exact-match** keyword to display only those routes with a set of communities exactly matching the list of specified communities. If you omit the **exact-match** keyword, those routes containing at least the specified communities are displayed.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp community** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp community 1820:1 exact-match

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 55
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16   192.168.40.24         0  1878  704  701  200  ?
*  10.16.0.0/16   192.168.40.24         0  1878  704  701  i
```

This table describes the significant fields shown in the display.

Table 11: show bgp community Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |

| Field | Description |
|-------------------|--|
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |

| Field | Description |
|--------|---|
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 28 | Creates an aggregate entry in a BGP routing table. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp , on page 268 | Displays BGP routes. |

show bgp convergence

To display whether a specific address family has reached convergence, use the **show bgp convergence** command in EXEC mode.

show bgp [ipv4 {unicast| multicast| labeled-unicast | all| tunnel}] convergence

show bgp [ipv6 {unicast}] convergence

show bgp [all {unicast| multicast| labeled-unicast | all| tunnel}] convergence

show bgp [vpn4 unicast] convergence

Syntax Description

| | |
|------------------------|--|
| ipv4 | (Optional) Specifies the IP Version 4 address family. |
| unicast | (Optional) Specifies the unicast address family. |
| multicast | (Optional) Specifies the multicast address family. |
| labeled-unicast | (Optional) Specifies unicast address prefixes. |
| all | (Optional) For subaddress family, specifies all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies the IP Version 6 address family. |
| all | (Optional) For address family, specifies all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Border Gateway Protocol (BGP) contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp convergence** command to see if there is any pending work for BGP to perform. The software checks the following conditions to determine whether the specified address family has converged. If all the conditions are true, the address family is considered converged.

- All received updates have been processed and best routes selected.
- All selected routes have been installed in the global Routing Information Base (RIB).
- All selected routes have been advertised to peers, including any peers that are not established (unless those peers have been administratively shut down). See the **shutdown (BGP)** command for more information about administrative shutdown.

While testing that all selected routes have been advertised to peers, the **show bgp convergence** command checks the size of the write queue for each neighbor. Because this queue is shared by all address families, there is a small possibility that the command indicates the address family has not converged when, in fact, it has converged. This could happen if the neighbor write queue contained messages from some other address family.

If the specified address family has not converged, the **show bgp convergence** command output does not indicate the amount of work that is pending. To display this information, use the **show bgp summary** command.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following shows the result of using the **show bgp convergence** command for an address family that has converged:

```
RP/0/RSP0/CPU0:router# show bgp convergence
Converged.
All received routes in RIB, all neighbors updated.
All neighbors have empty write queues.
```

The following shows the result of using the **show bgp convergence** command for an address family that has not converged:

```
RP/0/RSP0/CPU0:router# show bgp convergence
```

Not converged.
 Received routes may not be entered in RIB.
 One or more neighbors may need updating.

This table describes the significant fields shown in the display.

Table 12: show bgp convergence Field Descriptions

| Field | Description |
|----------------------------------|--|
| Converged/Not converged | Specifies whether or not all routes have been installed in the RIB and updates have been generated and sent to all neighbors. |
| [All] Received routes... | For convergence, all routes must have been installed into the RIB and all updates must have been generated. For non-convergence, some routes may not be installed in the RIB, or some routes that have been withdrawn have not yet been removed from the RIB, or some routes that are up to date in the RIB have not been advertised to all neighbors. |
| [All One or more] neighbors... | Specifies the status of neighbor updating. |

Related Commands

| Command | Description |
|--|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp summary , on page 442 | Displays the status of BGP peer connections. |
| shutdown (BGP) , on page 473 | Disables a neighbor without removing all of its configuration. |

show bgp dampened-paths

To display Border Gateway Protocol (BGP) dampened routes, use the **show bgp dampened-paths** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast | all}] dampened-paths
```

```
show bgp [ipv6 {unicast}] dampened-paths
```

```
show bgp [all {unicast| multicast| labeled-unicast | all| tunnel}] dampened-paths
```

```
show bgp [vpnv4 unicast [rd rd-address]] dampened-paths
```

```
show bgp [vrf {vrf-name| all} [ipv4 {unicast| labeled-unicast} | ipv6 unicast]] dampened-paths
```

Syntax Description

| | |
|--|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 {unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples The following is sample output from the **show bgp dampened-paths** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp dampened-paths

BGP router identifier 10.2.0.1, local AS number 3
BGP generic scan interval 60 secs
BGP main routing table version 7
Dampening enabled
BGP scan interval 60 secs
Status codes:s suppressed, d damped, h history, * valid, > best
                i - internal, S stale

Origin codes:i - IGP, e - EGP, ? - incomplete
  Network      From           Reuse      Path
*d 10.0.0.0    10.0.101.35    00:01:20  35 i
```

This table describes the significant fields shown in the display.

Table 13: show bgp dampened-paths Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|--------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network. |
| From | Neighbor from which the route was received. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 28 | Creates an aggregate entry in a BGP routing table. |
| bgp dampening , on page 76 | Enables BGP route dampening or changes various BGP route dampening factors. |
| clear bgp dampening , on page 121 | Clears BGP route dampening information and unsuppresses the suppressed routes. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp flap-statistics , on page 326 | Displays BGP routes that have flapped. |
| show bgp neighbors , on page 353 | Displays information about the TCP and BGP connections to neighbors. |

show bgp flap-statistics

To display information about Border Gateway Protocol (BGP) paths that have flapped, use the **show bgp flap-statistics** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast | all}] flap-statistics [regexp
regular-expression|route-policy route-policy-name |cidr-only |{ip-address| {mask| /prefix-length}}]
[longer-prefixes] [detail]
```

```
show bgp [ipv6 {unicast}] flap-statistics [regexp regular-expression|route-policy route-policy-name
|cidr-only |{ip-address| {mask| /prefix-length}}] [longer-prefixes] [detail]
```

```
show bgp [all {unicast| multicast| labeled-unicast | all}] flap-statistics [regexp
regular-expression|route-policy route-policy-name |cidr-only |{ip-address| {mask| /prefix-length}}]
[longer-prefixes] [detail]
```

```
show bgp [vpn4 unicast [rd rd-address]] flap-statistics [regexp regular-expression|route-policy
route-policy-name |cidr-only |{ip-address| {mask| /prefix-length}}] [longer-prefixes] [detail]
```

```
show bgp [vrf {vrf-name| all} [ipv4 | {unicast| labeled-unicast}| ipv6 unicast]] flap-statistics [regexp
regular-expression|route-policy route-policy-name |cidr-only |{ip-address| {mask| /prefix-length}}]
[longer-prefixes] [detail]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| vrf-name | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |

| | |
|---|--|
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| regex <i>regular-expression</i> | (Optional) Displays flap statistics for all paths that match the regular expression. |
| route-policy <i>route-policy-name</i> | (Optional) Displays flap statistics for a route policy. |
| cidr-only | (Optional) Displays only routes whose prefix length does not match the classful prefix length for that network. The cidr-only keyword can be specified only if the address family is IPv4. |
| <i>ip-address</i> | (Optional) Flap statistics for a network address only. |
| <i>mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| <i>/prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| longer-prefixes | (Optional) Displays flap statistics for the specified prefix and more-specific prefixes. The longer-prefixes keyword is available when the <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |
| detail | (Optional) Displays dampening parameters for the path. The detail keyword cannot be specified if the longer-prefixes keyword is specified. The detail keyword is available when the <i>ip-address</i> argument or <i>ip-address</i> and <i>mask</i> or <i>/prefix-length</i> arguments are specified. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Flap statistics are maintained only for paths if dampening is enabled using the **bgp dampening** command. If dampening is not enabled, the **show bgp flap-statistics** command does not display any paths.

If no arguments or keywords are specified, the software displays flap statistics for all paths for the specified address family. You can use the **regexp**, **filter-list**, **cidr-only**, and **longer-prefixes** options to limit the set of paths displayed.

If you specify a network address without a mask or prefix length, the longest matching prefix for the network address is displayed. When displaying flap statistics for a single route, use the **detail** keyword to display dampening parameters for the route.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp flap-statistics** command:

```
RP/0/RSP0/CPU0:router# show bgp flap-statistics

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 26180
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        From           Flaps Duration Reuse    Path
*d 10.0.0.0     172.20.16.177  4      00:13:31 00:18:10 100
*d 10.10.0.0    172.20.16.177  4      00:02:45 00:28:20 100
```

The following is sample output from the **show bgp flap-statistics** command with the **detail** keyword in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp flap-statistics 172.31.12.166 detail

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 738
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        From           Flaps Duration Reuse    Path
```

```

h 172.31.12.166          10.0.101.1      6      00:03:28      2 2000 3000
  Half life             Suppress        Reuse penalty   Max. supp. time
  00:15:00             2000           750            01:00:00

```

This table describes the significant fields shown in the display.

Table 14: show bgp flap-statistics Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP route identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|----------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network that is dampened. |
| From | IP address of the peer that advertised this route. |
| Flaps | Number of times the route has flapped. |
| Duration | Time (in hours:minutes:seconds) since the first flap. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path of the route that is being dampened. |
| Half life | Half-life value used when dampening this route. The half-life is the amount of time that must elapse to reduce the reuse penalty by half. The half-life value is specified using the bgp dampening command. |
| Suppress | Suppress value used to dampen this route. The suppress value is the value that the penalty must exceed for the route to be suppressed. The suppress value can be configured using the bgp dampening command. |
| Reuse penalty | Reuse penalty used to dampen this route. The penalty must fall below the reuse penalty for the route to be unsuppressed. The reuse penalty can be configured using the bgp dampening command. |
| Max supp. time | Maximum length of time that the route may be suppressed due to dampening. The maximum suppress time can be configured using the bgp dampening command. |

Related Commands

| Command | Description |
|---|---|
| bgp dampening , on page 76 | Enables BGP route dampening or changes various BGP route dampening factors. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp dampened-paths , on page 321 | Displays the BGP dampened routes. |
| show bgp neighbors , on page 353 | Displays information about BGP neighbors. |

show bgp inconsistent-as

To display Border Gateway Protocol (BGP) routes originated from more than one autonomous system, use the **show bgp inconsistent-as** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast | all| tunnel}] inconsistent-as [standby]
```

```
show bgp [ipv6 {unicast}] inconsistent-as [standby]
```

```
show bgp [all {unicast| multicast| labeled-unicast | all| tunnel}] inconsistent-as [standby]
```

```
show bgp vpnv4 unicast [rd rd-address] inconsistent-as [standby]
```

```
show bgp [vrf {vrf-name| all}] [ipv4 | {unicast| labeled-unicast}| ipv6 unicast]] inconsistent-as [standby]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or for the subaddress family, each matching routing table is examined in turn.

Use the **show bgp inconsistent-as** command to search through all prefixes in the specified BGP routing table and display the paths for any prefix that has inconsistent originating autonomous system numbers. The originating autonomous system is the last autonomous system number displayed in the path field and should be the same for all paths.

If a prefix has one or more paths originating from different autonomous systems, all paths for that prefix are displayed.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples The following is sample output from the **show bgp inconsistent-as** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp inconsistent-as

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 1129
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric          LocPrf Weight Path
```

show bgp inconsistent-as

```

* 10.0.0.0          172.16.232.55      0          0 300 88 90 99 ?
*>                172.16.232.52      2222       0 400 ?
* 172.16.0.0       172.16.232.55      0          0 300 90 99 88 200 ?
*>                172.16.232.52      2222       0 400 ?
* 192.168.199.0    172.16.232.55      0          0 300 88 90 99 ?
*>                172.16.232.52      2222       0 400 ?

```

This table describes the significant fields shown in the display.

Table 15: show bgp inconsistent-as Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|--------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|---|---|
| network (BGP), on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP), on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |

| Command | Description |
|-------------------|---|
| set default -safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp labels

To display Border Gateway Protocol (BGP) routes and their incoming and outgoing labels, use the **show bgp labels** command in EXEC mode.

show bgp labels

Syntax Description

| | |
|--|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled-unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd <i>rd-address</i> | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 {unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp labels** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp vrf BAR ipv4 unicast labels

BGP VRF BAR, state: Active BGP Route Distinguisher: 100:1 BGP router identifier 10.1.1.1,
local AS number 100
BGP table state: Active BGP main routing table version 12

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop        Rcvd Label      Local Label
Route Distinguisher: 100:1 (default for vrf BAR)
*> 20.1.1.1/32      10.0.101.1      16              no-label
*> 20.1.1.2/32      10.0.101.1      16              no-label
*> 20.1.1.3/32      10.0.101.1      16              no-label
*> 20.1.1.4/32      10.0.101.1      16              no-label
*> 20.1.1.5/32      10.0.101.1      16              no-label

Processed 5 prefixes, 5 paths
```

This table describes the significant fields shown in the display.

Table 16: show bgp labels Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP Route Distinguisher | BGP route distinguisher. |
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP table state | State of the BGP routing table. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Rcvd Label | Received label. |
| Local Label | Local label. |

Related Commands

| Command | Description |
|-------------------------|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default subaddress Family Identifier (SAFI) for the current session. |

show bgp l2vpn rd

To display routes with a specific route distinguisher under L2VPN address family, use the **show bgp l2vpn rd** command in EXEC mode.

```
show bgp l2vpn {vpls|vpws} rd rd_value
```

Syntax Description

| | |
|-----------------|--|
| vpls | Specifies Virtual Private LAN Services (VPLS). |
| vpws | Specifies Virtual Private Wire Service (VPWS). |
| <i>rd_value</i> | Value of the route distinguisher. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

Examples

The following example is sample output from the **show bgp l2vpn vpls** for route distinguisher 1:1:

```
RP/0/RSP0/CPU0:router#show bgp l2vpn vpls rd 1:1 2:1
BGP routing table entry for 2:1/32, Route Distinguisher: 1:1
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          2          2
  Local Label: 16000
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    100.100.100.1
```

```

Path #1: Received by speaker 0
Local
  0.0.0.0 from 0.0.0.0 (200.200.200.1)
  Origin IGP, localpref 100, valid, redistributed, best, import-candidate
  Extended community: RT:4:4 L2VPN:19:0:1500
  Block Size:10

```

The following example is sample output from the **show bgp l2vpn vpws** for route distinguisher 200:200:

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpws rd 200:200 3:1
BGP routing table entry for 3:1/32, Route Distinguisher: 200:200
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          6         6
  Local Label: 16015
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.1
  Advertised to peers (in unique update groups):
    100.100.100.1
Path #1: Received by speaker 0
Local
  0.0.0.0 from 0.0.0.0 (200.200.200.1)
  Origin IGP, localpref 100, valid, redistributed, best, import-candidate
  Extended community: RT:2:2 L2VPN:4:0:1500
  Circuit Vector:0xfd 0xff
  Block Size:10

```

The following example is sample output from the **show bgp l2vpn vpls** for local NLRI:

RD is 3.3.3.3:3276, NH Address is 100.0.0.1, and VPLS ID is 150:200. The RT is 200:100.


Note

The RT and VPLS-ID are always same for the same VPLS instance.

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls
Sat Jun  6 17:01:18.610 PST
BGP router identifier 3.3.3.3, local AS number 101
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0
BGP main routing table version 5
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
  Network      Next Hop
Route Distinguisher: 3.3.3.3:3276 (default for vrf gl:b1)
*>i200.0.0.1    30.0.0.2
*>i100.0.0.1    0.0.0.0
Route Distinguisher: 2.2.2.2:3435
*>i200.0.0.1    30.0.0.2

Processed 3 prefixes, 3 paths

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls rd 3.3.3.3:3276 100.0.0.1
Sat Jun  6 16:40:03.191 PST
BGP routing table entry for 100.0.0.1, Route Distinguisher: 3.3.3.3:3276
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          3         3
  Last Modified: Jun  6 11:20:57.944 for 05:19:05
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    30.0.0.2
Path #1: Received by speaker 0
Local
  0.0.0.0 from 0.0.0.0 (3.3.3.3)
  Origin IGP, localpref 100, valid, redistributed, best, import-candidate

```

```
Extended community: RT:200:100 VPLS-ID:150:200
```

The following example is sample output from the show bgp l2vpn vpls for remote NLRI:

RD is 2.2.2.2:3435, NH Address is 200.0.0.1, and VPLS ID is 150:200. The RT is 200:100.

```
RP/0/RSP0/CPU0:router#show bgp l2vpn vpls rd 2.2.2.2:3435 200.0.0.1
```

```
Sat Jun 6 16:53:55.726 PST
```

```
BGP routing table entry for 200.0.0.1, Route Distinguisher: 2.2.2.2:3435
```

```
Versions:
```

```
Process          bRIB/RIB  SendTblVer
Speaker          5         5
```

```
Last Modified: Jun 6 11:20:57.944 for 05:32:58
```

```
Paths: (1 available, best #1)
```

```
Not advertised to any peer
```

```
Path #1: Received by speaker 0
```

```
Local
```

```
30.0.0.2 from 30.0.0.2 (133.133.133.133)
```

```
Origin IGP, localpref 100, valid, internal, best, import-candidate, imported
```

```
Extended community: RT:200:50 VPLS-ID:150:200
```

show bgp l2vpn vpls

To display L2VPN information on BGP summary, routes for a specified bridge group domain, advertised routes, routes with a specific route distinguisher, BGP neighbor connections, nexthops, and BGP process, use the **show bgp l2vpn vpls** command in EXEC mode.

show bgp l2vpn vpls {summary | rd | neighbors | nexthops | bdomain | advertised | process}

Syntax Description

| | |
|-------------------|--|
| summary | Displays the summary of BGP neighbor status. |
| rd | Displays routes with a specific route distinguisher. |
| neighbors | Displays detailed information on TCP and BGP neighbor connections. |
| nexthops | Shows nexthop related information. |
| bdomain | Displays routes for a specified Bridge Group:domain |
| advertised | Shows advertised routes. |
| process | Displays BGP process information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |

Examples

The following example is sample output from the **show bgp l2vpn vpls** command:

```
RP/0/RSP0/CPU0:router#show bgp l2vpn vpls
Wed Mar 17 15:26:29.433 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop        Rcvd Label        Local Label
Route Distinguisher: 101:1 (default for vrf bg1:bg1_bd1)
*>i10.10.10.10/32    10.10.10.10    nolabel           nolabel
*> 60.60.60.60/32    0.0.0.0        nolabel           nolabel
Route Distinguisher: 102:1 (default for vrf bg1:bg1_bd2)
*>i10.10.10.10/32    10.10.10.10    nolabel           nolabel
*> 60.60.60.60/32    0.0.0.0        nolabel           nolabel
```

The following example is sample output from the **show bgp l2vpn vpls** command with the summary keyword:

```
RP/0/RSP0/CPU0:router#show bgp l2vpn vpls summary
Wed Mar 17 15:27:09.502 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process          RcvTblVer    bRIB/RIB    LabelVer    ImportVer    SendTblVer    StandbyVer
Speaker          24001        24001        24001        24001        24001        0

Neighbor         Spk          AS MsgRcvd  MsgSent     TblVer      InQ  OutQ  Up/Down    St/PfxRcd
10.10.10.10      0            1  45532      8392        24001     0    0  03:06:25    8000
```

The following example is sample output from the **show bgp l2vpn vpls** command for Route Distinguisher: 101:1:

```
RP/0/RSP0/CPU0:router#show bgp l2vpn vpls rd 101:1
Wed Mar 17 15:27:31.347 EDT
BGP router identifier 60.60.60.60, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0
BGP main routing table version 24001
BGP NSR Initial initsync version 1 (Reached)
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop        Rcvd Label        Local Label
Route Distinguisher: 101:1 (default for vrf bg1:bg1_bd1)
*>i10.10.10.10/32    10.10.10.10    nolabel           nolabel
*> 60.60.60.60/32    0.0.0.0        nolabel           nolabel
```

Processed 2 prefixes, 2 paths

The following example is sample output from the **show bgp l2vpn vpls** command for BGP neighbor 10.10.10.10:

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls neighbors 10.10.10.10
Wed Mar 17 15:28:28.766 EDT

BGP neighbor is 10.10.10.10
Remote AS 1, local AS 1, internal link
Remote router ID 10.10.10.10
BGP state = Established, up for 03:07:44
NSR State: None
Last read 00:00:31, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:45, attempted 19, written 19
Second last write 00:01:45, attempted 19, written 19
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Mar 17 15:27:57.362 last full not set pulse count 847
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Non-stop routing is enabled
Graceful restart is enabled
Restart time is 300 seconds
Stale path timeout time is 1200 seconds
Neighbor capabilities:
  Route refresh: advertised and received
  Graceful Restart (GR Awareness): received
  4-byte AS: advertised and received
  Address family IPv4 Unicast: advertised and received
  Address family VPNv4 Unicast: advertised and received
  Address family L2VPN VPLS: advertised and received
Received 45533 messages, 0 notifications, 0 in queue
Sent 8393 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 0 secs

For Address Family: IPv4 Unicast
BGP neighbor version 1
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
  Local restart time is 300, RIB purge time is 900 seconds
  Maximum stalepath time is 1200 seconds
  Remote Restart time is 300 seconds
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 524288
Threshold for warning message 75%, restart interval 0 min
An EoR was received during read-only mode
Last ack version 1, Last synced ack version 0
Outstanding version objects: current 0, max 0

For Address Family: VPNv4 Unicast
BGP neighbor version 1
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
  Local restart time is 300, RIB purge time is 900 seconds
  Maximum stalepath time is 1200 seconds
  Remote Restart time is 300 seconds
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths

```

```

Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 524288
Threshold for warning message 75%, restart interval 0 min
An EoR was received during read-only mode
Last ack version 1, Last synced ack version 0
Outstanding version objects: current 0, max 0

For Address Family: L2VPN VPLS
BGP neighbor version 24001
Update group: 0.2
AF-dependant capabilities:
  Graceful Restart Capability advertised and received
    Local restart time is 300, RIB purge time is 900 seconds
    Maximum stalepath time is 1200 seconds
    Remote Restart time is 300 seconds
Route refresh request: received 0, sent 203
8000 accepted prefixes, 8000 are bestpaths
Cumulative no. of prefixes denied: 18172.
  No policy: 0, Failed RT match: 18172
  By ORF policy: 0, By policy: 0
Prefix advertised 8000, suppressed 0, withdrawn 0
Maximum prefixes allowed 524288
Threshold for warning message 75%, restart interval 0 min
An EoR was received during read-only mode
Last ack version 24001, Last synced ack version 0
Outstanding version objects: current 0, max 2

Connections established 1; dropped 0
Local host: 60.60.60.60, Local port: 179
Foreign host: 10.10.10.10, Foreign port: 50472
Last reset 00:00:00

```

The following example is sample output from the **show bgp l2vpn vpls** command with the advertised keyword:

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls advertised
Wed Mar 17 15:29:13.787 EDT
Route Distinguisher: 101:1
60.60.60.60/32 is advertised to 10.10.10.10
  Path info:
    neighbor: Local          neighbor router id: 60.60.60.60
    valid redistributed best import-candidate
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    EXTCOMM
    origin: IGP
    aspath:
    extended community: RT:101:1 L2VPN AGI:1:101

```

The following example is sample output from the **show bgp l2vpn vpls** command with the nexthops keyword:

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls nexthops
Wed Mar 17 15:29:36.357 EDT
Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y46w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 03:10:50
  Time Spent: 0.000 secs

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000000
Nexthop Count: 2
Critical Trigger Delay: 3000msec
Non-critical Trigger Delay: 10000msec

```

show bgp l2vpn vpls

```

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.10.10.10   [R][NC][NL]  2           1/0       03:10:50 (Cri) 8000/8003

```

The following example is sample output from the show bgp l2vpn vpls command with the process keyword:

```

RP/0/RSP0/CPU0:router#show bgp l2vpn vpls process
Wed Mar 17 15:29:56.086 EDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASPLAIN
Autonomous System: 1
Router ID: 60.60.60.60 (manually configured)
Default Cluster ID: 60.60.60.60
Active Cluster IDs: 60.60.60.60
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 1200
RIB purge timeout time: 900
Non-stop routing is enabled
Update delay: 600
Generic scan interval: 60

Address family: L2VPN VPLS
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 24001
Table version synced to RIB: 1
RIB has not converged: version 0

Node          Process      Nbrs Estb Rst Upd-Rcvd Upd-Sent Nfn-Rcv Nfn-Snt
node0_RSP0_CPU0  Speaker      1     1   2   45347    237      0       0

```

show bgp neighbor-group

To display information about the Border Gateway Protocol (BGP) configuration for neighbor groups, use the **show bgp neighbor-group** command in EXEC mode.

```
show bgp neighbor-group group-name {configuration [defaults] [nvgen]| inheritance| users}
```

Syntax Description

| | |
|----------------------|--|
| <i>group-name</i> | Name of the address family group to display. |
| configuration | (Optional) Displays the effective configuration for the neighbor group, including any configuration inherited by this neighbor group. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in show running-config command output. If the defaults keyword is also specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | Displays the af-groups, session groups, and neighbor groups from which this neighbor group inherits configuration. |
| users | Displays the neighbors and neighbor groups that inherit configuration from this neighbor group. |

Command Default

No default behavior or value

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp neighbor-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a neighbor group, including any configuration inherited from session groups, address family groups, and neighbor groups through application of the **use** command. The source of each configured command is also displayed.

Use the **defaults** keyword to display all configuration for the neighbor group, including default configuration. The command output identifies default onfiguration. Use the **nvgen** keyword to display configuration in the output form of **show running-config** command. Output in this form is suitable for cutting and pasting into a configuration session.

The **show bgp neighbor-group** command with the *group-name* **inheritance** argument and keyword displays the session groups, address family groups, and neighbor groups from which the specified neighbor group inherits configuration.

The **show bgp neighbor-group** *group-name* command displays the neighbors and neighbor groups that inherit configuration from the specified neighbor group.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The examples use the following configuration:

```
af-group group3 address-family ipv4 unicast
  remove-private-AS
  soft-reconfiguration inbound
!
af-group group2 address-family ipv4 unicast
  use af-group group3
  send-community-ebgp
  send-extended-community-ebgp
  capability orf prefix both
!
session-group group3
  dmzlink-bw
!
neighbor-group group3
  use session-group group3
  timers 30 90
!
neighbor-group group1
  remote-as 1982
  use neighbor-group group2
  address-family ipv4 unicast
!
!
neighbor-group group2
  use neighbor-group group3
  address-family ipv4 unicast
  use af-group group2
  weight 100
!
```

The following is sample output from the **show bgp neighbor-group** command with the **configuration** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbor-group group1 configuration

neighbor-group group1
  remote-as 1982                []
  timers 30 90                 [n:group2 n:group3]
  dmzlink-bw                   [n:group2 n:group3 s:group3]
  address-family ipv4 unicast  []
  capability orf prefix both   [n:group2 a:group2]
  remove-private-AS           [n:group2 a:group2 a:group3]
  send-community-ebgp         [n:group2 a:group2]
  send-extended-community-ebgp [n:group2 a:group2]
```

```

soft-reconfiguration inbound    [n:group2 a:group2 a:group3]
weight 100                     [n:group2]
    
```

The configuration source is shown to the right of each command. In the output, the **remote-as** command is configured directly on neighbor group group1, and the **send-community-ebgp** command is inherited from neighbor group group2, which in turn inherits the setting from af-group group2.

The following is sample output from the **show bgp neighbor-group** command with the **users** keyword. This output shows that the group1 neighbor group inherits session (address family-independent configuration parameters) from the group2 neighbor group. The group1 neighbor group also inherits IPv4 unicast configuration parameters from the group2 neighbor group:

```

RP/0/RSP0/CPU0:router# show bgp neighbor-group group2 users

Session:      n:group1
IPv4 Unicast: n:group1
    
```

The following is sample output from the **show bgp neighbor-group** command with the **inheritance** keyword. This output shows that the specified neighbor group group1 inherits session (address family-independent configuration) from neighbor group group2, which inherits its own session from neighbor group group3. Neighbor group group3 inherited its session from session group group3. It also shows that the group1 neighbor-group inherits IPv4 unicast configuration parameters from the group2 neighbor group, which in turn inherits them from the group2 af-group, which itself inherits them from the group3 af-group:

```

RP/0/RSP0/CPU0:router# show bgp neighbor-group group1 inheritance

Session:      n:group2 n:group3 s:group3
IPv4 Unicast: n:group2 a:group2 a:group3
    
```

This table describes the significant fields shown in the display.

Table 17: show bgp neighbor-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified address family group. |
| s: | Indicates the name that follows is a session group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the setting is not explicitly configured or inherited, and the default value for the setting is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the default is for the setting to be disabled. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|--|
| af-group , on page 26 | Configures a BGP address family group. |
| session-group , on page 264 | Creates a session group and enters session group configuration mode. |
| show bgp af-group , on page 301 | Displays information about configuration for address family groups. |
| show bgp neighbors , on page 353 | Displays information about BGP neighbors, including configuration inherited from neighbor groups, session groups, and address family groups. |
| show bgp session-group , on page 436 | Displays information about the BGP configuration for session groups. |
| show running-config | Displays the contents of the currently running configuration or a subset of that configuration. |
| use , on page 514 | Inherits configuration from a neighbor group, a session group, or an address family group. |

show bgp neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp neighbors** command in EXEC mode.

show bgp neighbors [**performance-statistics**| **missing-eor**] [**standby**]

show bgp neighbors *ip-address* [**advertised-routes**| **dampened-routes**| **flap-statistics**| **performance-statistics**| **received** | {**prefix-filter**| **routes**}| **routes**] [**standby**]

show bgp neighbors *ip-address* [**configuration**| [**defaults**] | **nvgen**| **inheritance**][**standby**]

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| performance-statistics | (Optional) Displays performance statistics relative to work done by the BGP process for this neighbor. |
| missing-eor | (Optional) Displays neighbors that did not send end-of-rib (EoR) notification in read-only mode. |
| <i>ip-address</i> | (Optional) IP address of the BGP-speaking neighbor. If you omit this argument, all neighbors are displayed. |

| | |
|--|---|
| advertised-routes | (Optional) Displays all routes the router advertised to the neighbor. |
| dampened-routes | (Optional) Displays the dampened routes that are learned from the neighbor. |
| flap-statistics | (Optional) Displays flap statistics of the routes learned from the neighbor. |
| received { prefix-filter routes } | (Optional) Displays information received from the BGP neighbor. The options are: prefix-filter — Displays the prefix list filter. routes —Displays routes from the neighbor before inbound policy |
| routes | (Optional) Displays routes learned from the neighbor. |
| configuration | (Optional) Displays the effective configuration for the neighbor, including any settings that have been inherited from session groups, neighbor groups, or af-groups used by this neighbor. |
| defaults | (Optional) Displays all configuration settings, including any default settings. |
| nvgen | (Optional) Displays output in the show running-config command output. |
| inheritance | (Optional) Displays the session groups, neighbor groups, and af-groups from which this neighbor inherits configuration settings. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.1.1 | The command output was modified to display from BGP Accept Own configuration. |
| Release 4.0.0 | The command output was modified to include information on BGP additional paths send and receive information. |
| Release 4.3.0 | The command output was modified to include information from update wait-install configuration. |

| Release | Modification |
|---------------|---|
| Release 5.1.1 | The command output was modified to display the status of permanent paths. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify which routing table should be examined. If the **all** keyword is specified for address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp neighbors** command to display detailed information about all neighbors or a specific neighbor. Use the **performance-statistics** keyword to display information about the work related to specific neighbors done by the BGP process.

Use the **show bgp neighbors** command with the *ip-address* **received prefix-filter** argument and keyword to display the Outbound Route Filter (ORF) received from a neighbor.

Use the **advertised-routes** keyword to display a summary of the routes advertised to the specified neighbor.

Use the **dampened-routes** keyword to display routes received from the specified neighbor that have been suppressed due to dampening. For more details, see the **show bgp dampened-paths** command.

To display information about flapping routes received from a neighbor, use the **flap-statistics** keyword. For more details, see the **show bgp flap-statistics** command.

To display the routes received from a neighbor, use the **routes** keyword. For more details, see the **show bgp** command.

Use the **show bgp neighbor** command with the *ip-address* **configuration** argument and keyword to display the effective configuration of a neighbor, including configuration inherited from session groups, neighbor groups, or af-groups through application of the **use** command. Use the **defaults** keyword to display the value of all configurations for the neighbor, including default configuration. Use the **nvgen** keyword to display configuration output format of the **show running-config** command. Output in this format is suitable for cutting and pasting into a configuration session. Use the **show bgp neighbors** command with the *ip-address* **inheritance** argument and keyword to display the session groups, neighbor groups, and af-groups from which the specified neighbor inherits configuration.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp neighbors** command:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.1

BGP neighbor is 10.0.101.1, remote AS 2, local AS 1, external link
  Description: routem neighbor
  Remote router ID 10.0.101.1
  BGP state = Established, up for 00:00:56
  TCP open mode: passive only
BGP neighbor is 1.1.1.2
  Remote AS 300, local AS 100, external link
  Remote router ID 0.0.0.0
  BGP state = Idle (LC/FIB for the neighbor in reloading)
  Last read 00:00:00, Last read before reset 00:05:12
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3

  BFD enabled (session initializing)
  Last read 00:00:55, hold time is 180, keepalive interval is 60 seconds
  DMZ-link bandwidth is 1000 Mb/s
  Neighbor capabilities:
    Route refresh: advertised
    4-byte AS: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Multicast: advertised and received
  Received 119 messages, 0 notifications, 0 in queue
  Sent 119 messages, 22 notifications, 0 in queue
  Minimum time between advertisement runs is 60 seconds

For Address Family: IPv4 Unicast
  BGP neighbor version 137
  Update group: 1.3
  Community attribute sent to this neighbor
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: advertised
      Receive-mode: advertised
    Route refresh request: received 0, sent 0
    Policy for incoming advertisements is pass-all
    Policy for outgoing advertisements is pass-all
    5 accepted prefixes, 5 are bestpaths
    Prefix advertised 3, suppressed 0, withdrawn 0, maximum limit 1000000
    Threshold for warning message 75%

For Address Family: IPv4 Multicast
  BGP neighbor version 23
  Update group: 1.2
  Route refresh request: received 0, sent 0
  Policy for incoming advertisements is pass-all
  Policy for outgoing advertisements is pass-all
  2 accepted prefixes, 2 are bestpaths
  Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 131072
  Threshold for warning message 75%

Connections established 9; dropped 8
Last reset 00:02:10, due to User clear requested (CEASE notification sent - administrative
reset)
Time since last notification sent to neighbor: 00:02:10
Error Code: administrative reset
Notification data sent:
  None
```

This table describes the significant fields shown in the display.

Table 18: show bgp neighbors Field Descriptions

| Field | Description |
|--------------------------------------|---|
| BGP neighbor | IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external. |
| Description | Neighbor specific description. |
| remote AS | <ul style="list-style-type: none"> • Number of the autonomous system to which the neighbor belongs. • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. |
| local AS | <p>Autonomous system number of the local system.</p> <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) is asdot format is 1.0 to 65535.65535. |
| internal link | Neighbor is an internal BGP peer. |
| external link | Neighbor is an external BGP peer. |
| Administratively shut down | Neighbor connection is disabled using the shutdown command. |
| remote router ID | Router ID (an IP address) of the neighbor. |
| Neighbor under common administration | Neighbor is internal or a confederation peer. |
| BGP state | Internal state of this BGP connection. |
| BFD enabled | Status of bidirectional forwarding detection. |

| Field | Description |
|-----------------------|--|
| TCP open mode | TCP mode used in establishing the BGP session. The following valid TCP mode are supported: <ul style="list-style-type: none"> • default—Accept active/passive connections • passive-only—Accept only passive connections • active-only—Accept only active connections initiated by the router |
| Last read | Time since BGP last read a message from this neighbor. |
| hold time | Hold time (in seconds) used on the connection with this neighbor. |
| keepalive interval | Interval for sending keepalives to this neighbor. |
| DMZ-link bandwidth | DMZ link bandwidth for this neighbor. |
| Neighbor capabilities | BGP capabilities advertised and received from this neighbor. The following valid BGP capabilities are supported: <ul style="list-style-type: none"> • Multi-protocol • Route refresh • Graceful restart • Outbound Route Filter (ORF) type (128) Prefix |
| Route refresh | Indicates that the neighbor supports dynamic soft reset using the route refresh capability. |
| 4-byte AS | Indicates that the neighbor supports the 4-byte AS capability. |
| Address family | Indicates that the local system supports the displayed address family capability. If "received" is displayed, the neighbor also supports the displayed address family. |
| Received | Number of messages received from this neighbor, the number of notification messages received and processed from this neighbor, and the number of messages that have been received, but not yet processed. |

| Field | Description |
|--|---|
| Sent | Number of messages sent to this neighbor, the number of notification messages generated to be sent to this neighbor, and the number of messages queued to be sent to this neighbor. |
| Minimum time between advertisement runs | Advertisement interval (in seconds) for this neighbor. |
| For Address Family | Information that follows is specific to the displayed address family. |
| BGP neighbor version | Last version of the BGP database that was sent to the neighbor for the specified address family. |
| Update group | Update group to which the neighbor belongs. |
| Route reflector client | Indicates that the local system is acting as a route reflector for this neighbor. |
| Inbound soft reconfiguration allowed | Indicates that soft reconfiguration is enabled for routes received from this neighbor. Note If the neighbor has route refresh capability, then soft configuration received-only routes are not stored by the local system unless “override route refresh” is displayed. |
| eBGP neighbor with no inbound or outbound policy: defaults to drop | Indicates that the neighbor does not have an inbound or outbound policy configured using the route-policy (BGP) command. Hence, no routes are accepted from or advertised to this neighbor. |
| Private AS number removed from updates to this neighbor | Indicates that remove-private-AS is configured on the specified address family for this neighbor. |
| NEXT_HOP is always this router | Indicates that next-hop-self is configured on the specified address family for this neighbor. |
| Community attribute sent to this neighbor | Indicates that send-community-ebgp is configured on the specified address family for this neighbor. |
| Extended community attribute sent to this neighbor | Indicates that send-extended-community-ebgp is configured on the specified address family for this neighbor. |
| Default information originate | Indicates that default-originate is configured on the specified address family for this neighbor, together with the policy used, if one was specified in the default-originate configuration. An indication of whether the default route has been advertised to the neighbor is also shown. |

| Field | Description |
|---|---|
| AF-dependant capabilities | <p>BGP capabilities that are specific to a particular address family. The following valid AF-dependent BGP capabilities are supported:</p> <ul style="list-style-type: none"> • route refresh capability • route refresh capability OLD value |
| Outbound Route Filter | <p>Neighbor has the Outbound Route Filter (ORF) capability for the specified address family. Details of the capabilities supported are also shown:</p> <p>Send-mode—“advertised” is shown if the local system can send an outbound route filter to the neighbor. “received” is shown if the neighbor can send an outbound route filter to the local system.</p> <p>Receive-mode—“advertised” is shown if the local system can receive an outbound route filter from the neighbor. “received” is shown if the neighbor can receive an outbound route filter from the local system.</p> |
| Graceful Restart Capability | Indicates whether graceful restart capability has been advertised to and received from the neighbor for the specified address family. |
| Neighbor preserved the forwarding state during latest restart | Indicates that when the neighbor connection was last established, the neighbor indicated that it preserved its forwarding state for the specified address family. |
| Local restart time | Restart time (in seconds) advertised to this neighbor. |
| RIB purge time | RIB purge time (in seconds) used for graceful restarts. |
| Maximum stalepath time | Maximum time (in seconds) a path received from this neighbor may be marked as stale if the neighbor restarts. |
| Remote Restart time | Restart time received from this neighbor. |
| Route refresh request | Number of route refresh requests sent and received from this neighbor. |

| Field | Description |
|---|--|
| Outbound Route Filter (ORF) | <p>“sent” indicates that an outbound route filter has been sent to this neighbor. “received” indicates that an outbound route filter has been received from this neighbor.</p> <p>Note A received outbound route filter may be displayed using the show bgp neighbors command with the received prefix-filter keywords.</p> |
| First update is deferred until ORF or ROUTE-REFRESH is received | If the local system advertised the receive capability and the neighbor has advertised send capability, no updates are generated until specifically asked by the neighbor (using a ROUTE-REFRESH or ORF with immediate request). |
| Scheduled to send the Prefix-list filter | Indicates the local system is due to send an outbound route filter request in order to receive updates from the neighbor. |
| Inbound path policy | Indicates if an inbound path policy is configured. |
| Outbound path policy | Indicates if an outbound path policy is configured. |
| Incoming update prefix filter list | Indicates a prefix list is configured to filter inbound updates from the neighbor. |
| Default weight | Default weight for routes received from the neighbor. |
| Policy for incoming advertisements | Indicates a route policy is configured to be applied to inbound updates from the neighbor. |
| Policy for outgoing advertisements | Indicates a route policy is configured to be applied to outbound updates to the neighbor. |
| Type | <p>Indicates whether the condition map selects routes that should be advertised, or routes that should not be advertised:</p> <p>Exist—Routes advertised if permitted by the condition route map.</p> <p>Non-exist—Routes advertised if denied by the condition route map.</p> |
| accepted prefixes | Number of prefixes accepted. |
| Prefix advertised | Number of prefixes advertised to the neighbor during the lifetime of the current connection with the neighbor. |

| Field | Description |
|---|--|
| suppressed | Number of prefix updates that were suppressed because no transitive attributes changed from one best path to the next. Note Update suppression occurs only for external BGP neighbors. |
| withdrawn | Number of prefixes withdrawn from the neighbor during the lifetime of the current connection with the neighbor. |
| maximum limit | Maximum number of prefixes that may be received from the neighbor. If "(warning-only)" is displayed, a warning message is generated when the limit is exceeded, otherwise the neighbor connection is shut down when the limit is exceeded. |
| Threshold for warning message | Percentage of maximum prefix limit for the neighbor at which a warning message is generated. |
| Connections established | Number of times the router has established a BGP peering session with the neighbor. |
| dropped | Number of times that a good connection has failed or been taken down. |
| Last reset due to | Reason that the connection with the neighbor was last reset. |
| Time since last notification sent to neighbor | Amount of time since a notification message was last sent to the neighbor. |
| Error Code | Type of notification that was sent. The notification data, if any, is also displayed. |
| Time since last notification received from neighbor | Amount of time since a notification message was last received from the neighbor. |
| Error Code | Type of notification that was received. The notification data received, if any, is also displayed |
| External BGP neighbor may be up to <n> hops away | Indicates ebgp-multihop is configured for the neighbor. |
| External BGP neighbor not directly connected | Indicates that the neighbor is not directly attached to the local system. |
| Notification data sent: | Data providing more details on the error along with the error notification sent to the neighbor. |

The following is sample output from the **show bgp neighbors** command with the **advertised-routes** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 172.20.16.178 routes

BGP router identifier 172.20.16.181, local AS number 1
BGP main routing table version 27
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        Next Hop          Metric LocPrf Weight Path
*> 10.0.0.0      172.20.16.178        40             0 10 ?
*> 10.22.0.0     172.20.16.178        40             0 10 ?
```

The following is sample output from the **show bgp neighbors** command with the **routes** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.1 dampened-routes

BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network        From          Reuse    Path
*d 10.0.0.0     10.0.101.1    00:59:30 2 100 1000 i
*d 11.0.0.0     10.0.101.1    00:59:30 2 100 1000 i
*d 12.0.0.0     10.0.101.1    00:59:30 2 100 1000 i
*d 13.0.0.0     10.0.101.1    00:59:30 2 100 1000 i
*d 14.0.0.0     10.0.101.1    00:59:30 2 100 1000 i
```

This table describes the significant fields shown in the display.

Table 19: show bgp neighbors routes Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |

| Field | Description |
|--------|---|
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

The following is sample output from the **show bgp neighbors** command with the **dampened-routes** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.1 flap-statistics
BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 48
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          From            Flaps Duration Reuse      Path
   -----          -
h 10.1.0.0          10.0.101.1      5008 2d02h          2 5000 1000
h 10.2.0.0          10.0.101.1      5008 2d02h          2 2000 3000
h 10.2.0.0          10.0.101.1      5008 2d02h          2 9000 6000
*d 10.0.0.0         10.0.101.1      5008 2d02h          00:59:30 2 100 1000
h 10.0.0.0/16      10.0.101.1      5008 2d02h          2 100 102
*d 10.11.0.0       10.0.101.1      5008 2d02h          00:59:30 2 100 1000
*d 10.12.0.0       10.0.101.1      5008 2d02h          00:59:30 2 100 1000
*d 10.13.0.0       10.0.101.1      5008 2d02h          00:59:30 2 100 1000
*d 10.14.0.0       10.0.101.1      5008 2d02h          00:59:30 2 100 1000
h 192.168.0.0/16  10.0.101.1      5008 2d02h          2 100 101
```

This table describes the significant fields shown in the display.

Table 20: show bgp neighbors dampened-routes Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |

| Field | Description |
|--------------|--|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>>—Path is the best path to use for that network.</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| From | Neighbor from which the route was received. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

The following is sample output from the **show bgp neighbors** command with the **flap-statistics** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.2 performance-statistics

BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)

  Received 3023 messages, 0 notifications, 0 in queue
  Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 21: show bgp neighbors flap-statistics Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP route identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between when the BGP process scans for the specified address family and subaddress family. |

| Field | Description |
|--------------|---|
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <p>S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned.</p> <p>s—Path is more specific than a locally sourced aggregate route and has been suppressed.</p> <p>*—Path is valid.</p> <p>The second character may be (in order of precedence):</p> <p>d—Path is dampened.</p> <p>h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid.</p> <p>The third character may be:</p> <p>i—Path was learned by an internal BGP (iBGP) session.</p> |
| Origin codes | <p>Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |
| Network | IP prefix and prefix length for a network. |
| From | IP address of the peer that advertised this route. |
| Flaps | Number of times the route has flapped. |
| Duration | Time (in hours:minutes:seconds) since the router noticed the first flap. |
| Reuse | Time (in hours:minutes:seconds) after which the path is made available. |
| Path | Autonomous system path to reach the destination network. |

The following is sample output from the **show bgp neighbors** command with the **performance-statistics** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.2 performance-statistics
BGP neighbor is 10.0.101.2, remote AS 1
  Read 3023 messages (58639 bytes) in 3019 calls (time spent: 1.312 secs)
  Read throttled 0 times
  Processed 3023 inbound messages (time spent: 0.198 secs)
  Wrote 58410 bytes in 6062 calls (time spent: 3.041 secs)
  Processing write list: wrote 0 messages in 0 calls (time spent: 0.000 secs)
  Processing write queue: wrote 3040 messages in 3040 calls (time spent: 0.055 secs)
  Received 3023 messages, 0 notifications, 0 in queue
  Sent 3040 messages, 0 notifications, 0 in queue
```

This table describes the significant fields shown in the display.

Table 22: show bgp neighbors performance-statistics Field Descriptions

| Field | Description |
|------------------------|--|
| Read | Indicates the number of messages received from the neighbor, the total size of received messages, the number of read operations performed, and the real time spent (in seconds) by the process performing read operations for this neighbor. |
| Read throttled | Number of times that reading from the TCP connection to this neighbor has been throttled. Throttling is due to a backlog of messages that have been read but not processed. |
| inbound messages | Number of read messages that have been processed, and the real time spent processing inbound messages for this neighbor. |
| Wrote | Amount of data that has been sent to this neighbor, number of write operations performed, and the real time spent by the process performing write operations for this neighbor. |
| Processing write list | Number of messages written from the write list to this neighbor, number of times the write list has been processed, and real time spent processing the write list. Note Write lists typically contain only update messages. |
| Processing write queue | Number of messages written from the write queue to this neighbor, number of times the write queue has been processed, and real time spent processing the write queue. |

| Field | Description |
|----------|---|
| Received | Number of messages received from this neighbor, number of notification messages received and processed from this neighbor, and number of messages that have been received, but not yet processed. |
| Sent | Number of messages sent to this neighbor, number of notification messages generated to be sent to this neighbor, and number of messages queued to be sent to this neighbor. |

The following is sample output from the **show bgp neighbors** command with the **configuration** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors 10.0.101.1 configuration
neighbor 10.0.101.1
  remote-as 2 []
  bfd fast-detect []
  address-family ipv4 unicast []
  policy pass-all in []
  policy pass-all out []
  address-family ipv4 multicast []
  policy pass-all in []
  policy pass-all out []
```

This table describes the significant fields shown in the display.

Table 23: show bgp neighbors configuration Field Descriptions

| Field | Description |
|---------------------------|--|
| neighbor | IP address configuration of the neighbor. |
| remote-as | Remote autonomous system configured on the neighbor. |
| bfd fast-detect | BFD parameter configured on the neighbor. |
| address-family | Address family and subsequent address family configured on the router. |
| route-policy pass-all in | Route policy configured for inbound updates. |
| route-policy pass-all out | Route policy configured for outbound updates. |

The following sample output shows sample output from **show bgp neighbors** command with additional paths send and receive capabilities advertised to neighbors:

```
BGP neighbor is 80.0.0.30
  Remote AS 100, local AS 100, internal link
  Remote router ID 33.33.33.33
```

```

BGP state = Established, up for 19:54:12
NSR State: None
Last read 00:00:25, Last read before reset 19:54:54
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:02, attempted 19, written 19
Second last write 00:01:02, attempted 19, written 19
Last write before reset 19:54:54, attempted 29, written 29
Second last write before reset 19:54:59, attempted 19, written 19
Last write pulse rcvd Nov 11 12:58:03.838 last full not set pulse count 2407
Last write pulse rcvd before reset 19:54:54
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 19:54:54, second last 19:54:54
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 19:54:54, second last 19:54:59
Precedence: internet
Non-stop routing is enabled
Graceful restart is enabled
Restart time is 120 seconds
Stale path timeout time is 360 seconds
Neighbor capabilities:
  Adv          Rcvd
Route refresh: Yes          Yes
4-byte AS:    Yes          Yes
Address family IPv4 Unicast: Yes          Yes
Address family IPv4 Labeled-unicast: Yes      Yes
Address family VPNv4 Unicast: Yes          Yes
Address family IPv6 Unicast: Yes          Yes
Address family VPNv6 Unicast: Yes          Yes
Address family IPv4 MDT:    Yes          Yes
Message stats:
  InQ depth: 0, OutQ depth: 0
  Last_Sent      Sent  Last_Rcvd      Rcvd
Open:           Nov 10 17:03:52.731    2    Nov 10 17:03:52.730    2
Notification:   ---          0    ---                0
Update:         Nov 10 17:05:02.435    20   Nov 10 17:04:58.812    12
Keepalive:      Nov 11 12:58:03.632    1197 Nov 11 12:57:40.458    1196
Route_Refresh:  ---          0    ---                0
Total:          1219          1210
Minimum time between advertisement runs is 0 secs

```

```

For Address Family: IPv4 Unicast
BGP neighbor version 13
Update group: 0.9
NEXT_HOP is always this router
AF-dependant capabilities:
  Graceful Restart capability advertised and received
  Neighbor preserved the forwarding state during latest restart
  Local restart time is 120, RIB purge time is 600 seconds
  Maximum stalepath time is 360 seconds
  Remote Restart time is 120 seconds
  Additional-paths Send: advertised and received
  Additional-paths Receive: advertised and received
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Prefix advertised 10, suppressed 0, withdrawn 0, maximum limit 524288
Threshold for warning message 75%
AIGP is enabled
An EoR was received during read-only mode
Last ack version 13, Last synced ack version 0
Outstanding version objects: current 0, max 1
Additional-paths operation: Send and Receive

```

```

For Address Family: IPv4 Labeled-unicast
BGP neighbor version 13
Update group: 0.4 (Update Generation Throttled)

```

```

AF-dependant capabilities:
  Graceful Restart capability advertised and received
  Neighbor preserved the forwarding state during latest restart
  Local restart time is 120, RIB purge time is 600 seconds
  Maximum stalepath time is 360 seconds
  Remote Restart time is 120 seconds

```

```

    Additional-paths Send: received
    Additional-paths Receive: received
    Route refresh request: received 0, sent 0
    0 accepted prefixes, 0 are bestpaths
    Prefix advertised 2, suppressed 0, withdrawn 0, maximum limit 131072
    Threshold for warning message 75%
    AIGP is enabled
    An EoR was received during read-only mode
    Last ack version 13, Last synced ack version 0
    Outstanding version objects: current 0, max 1
    Additional-paths operation: None

```

This is sample output of the **show bgp neighbors** command when update wait-install is enabled. If the session open is postponed due to the reloading of the LC/FIB, the text "LC/FIB for the neighbor in reloading" is displayed next to the BGP state.

```

RP/0/RSP0/CPU0:router#show bgp neighbors 1.1.1.2

BGP neighbor is 1.1.1.2
Remote AS 300, local AS 100, external link
Remote router ID 0.0.0.0
BGP state = Idle (LC/FIB for the neighbor in reloading)
Last read 00:00:00, Last read before reset 00:05:12
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3

```

This is sample output from **show bgp neighbors** command that displays status of Accept Own configuration:

```

RP/0/RSP0/CPU0:router#show bgp neighbors 45.1.1.1

BGP neighbor is 45.1.1.1
Remote AS 100, local AS 100, internal link
Remote router ID 45.1.1.1
BGP state = Established, up for 00:19:54
NSR State: None
Last read 00:00:55, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:54, attempted 19, written 19
Second last write 00:01:54, attempted 19, written 19
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Jul 19 11:45:38.776 last full not set pulse count 43
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Non-stop routing is enabled
Neighbor capabilities:
  Route refresh: advertised and received
  4-byte AS: advertised and received
  Address family VPNv4 Unicast: advertised and received
  Address family VPNv6 Unicast: advertised and received
Received 22 messages, 0 notifications, 0 in queue
Sent 22 messages, 0 notifications, 0 in queue
Minimum time between advertisement runs is 0 secs

For Address Family: VPNv4 Unicast

BGP neighbor version 549
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is drop_111.x.x.x
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 524288

```

```

Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
Accept-own is enabled
An EOR was received during read-only mode
Last ack version 549, Last synced ack version 0
Outstanding version objects: current 0, max 0
Additional-paths operation: None

For Address Family: VPNv6 Unicast

BGP neighbor version 549
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 0
Policy for incoming advertisements is pass-all
Policy for outgoing advertisements is drop_111.x.x.x
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 524288
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
Accept-own is enabled
An EOR was received during read-only mode
Last ack version 549, Last synced ack version 0
Outstanding version objects: current 0, max 0
Additional-paths operation: None

Connections established 1; dropped 0
Local host: 15.1.1.1, Local port: 179
Foreign host: 45.1.1.1, Foreign port: 56391
Last reset 00:00:00
RP/0/0/CPU0:BGP1-6#

```

This sample output from the **show bgp neighbor** command displays the status of permanent paths:

```

RP/0/RSP0/CPU0:router#show bgp neighbors 3.3.3.3
BGP neighbor is 3.3.3.3
Remote AS 30813, local AS 30813, internal link
Remote router ID 3.3.3.3
BGP state = Established, up for 01:39:14
Last read 00:00:58, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:53, attempted 2054, written 2054
Second last write 00:00:53, attempted 45, written 45
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Aug 14 07:53:56.846 last full not set pulse count 226
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Multi-protocol capability received
Neighbor capabilities:
  Route refresh:          Adv      Rcvd
                        Yes      Yes
  4-byte AS:             Yes      Yes
  Address family IPv4 Unicast: Yes      Yes

For Address Family: IPv4 Unicast
BGP neighbor version 1111
Update group: 0.3 Filter-group: 0.5 No Refresh request being processed
NEXT HOP is always this router
Default information originate: default sent
AF-dependent capabilities:
  Additional-paths Send: received
  Additional-paths Receive: received
Route refresh request: received 0, sent 0
Policy for incoming advertisements is PASS
Policy for outgoing advertisements is PASS

```

show bgp neighbors

```

100 accepted prefixes, 100 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 5500, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 0, Last synced ack version 0
Outstanding version objects: current 1, max 1
Additional-paths operation: None
Advertise Permanent-Network enabled

Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 64742
Last reset 00:00:00

```

Related Commands

| Command | Description |
|--|---|
| clear bgp , on page 119 | Resets a BGP connection or session. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp , on page 268 | Displays entries in the BGP routing table. |
| show bgp dampened-paths , on page 321 | Displays BGP dampened routes. |
| show bgp flap-statistics , on page 326 | Displays BGP routes that have flapped. |
| show bgp neighbor-group , on page 349 | Displays information about the BGP configuration for neighbor groups. |
| shutdown (BGP) , on page 473 | Disables a neighbor without removing all of its configuration. |

show bgp neighbors nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information across neighbors, use the **show bgp neighbors nsr** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| all}| ipv6 {unicast| multicast| all}| vpnv4 unicast| vpnv6 unicast| vrf
{all| vrf_name}] neighbors nsr [standby]
```

Syntax Description

| | |
|----------------------|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vpnv6 unicast | (Optional) Specifies VPNv6 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf_name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| standby | (Optional) Displays information about the standby card. |

Command Default

No default behavior or values.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp neighbors nsr** command with the **standby** keyword:

```
RP/0/RSP0/CPU0:router# show bgp neighbors nsr standby

BGP neighbor is 2.2.2.2
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0

BGP neighbor is 10.0.101.5
  BGP state = Established, up for 05:19:00
  NSR state = NSR Ready
  Outstanding Postits: 0

BGP neighbor is 10.1.0.5
  BGP state = Established, up for 5d04h
  NSR state = NSR Ready
  Outstanding Postits: 0
```

This table describes the significant fields shown in the display.

Table 24: show bgp neighbors nsr Field Descriptions

| Field | Description |
|---------------------|---|
| BGP state | Displays BGP neighbor peering state. |
| NSR state | Displays BGP neighbor NSR state. |
| Outstanding Postits | Displays the postit counters of pending events. |

Related Commands

| Command | Description |
|---|---|
| nsr (BGP), on page 210 | Activates the Border Gateway Protocol (BGP) nonstop routing (NSR). |
| show bgp summary nsr, on page 447 | Displays the Border Gateway Protocol (BGP) nonstop routing (NSR) information. |
| show bgp summary, on page 442 | Displays the status of all Border Gateway Protocol (BGP) connections. |

show bgp nexthops

To display statistical information about the Border Gateway Protocol (BGP) next hops, use the **show bgp nexthops** command in EXEC mode.

```
show bgp nexthops [statistics] [speaker speaker-id] []
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled-unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| statistics | (Optional) Specifies nexthop statistics. |
| speaker <i>speaker-id</i> | (Optional) Specifies a speaker process ID. |

Command Default

No default behavior or value

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show bgp nexthops** command displays statistical information about next-hop notifications, the time spent processing the notifications, and details about each next-hop that has been registered with the Routing Information Base (RIB).

Use the **vrf** *vrf-name* keyword and argument to display only the next-hops present in the specified VPN routing and forwarding (VRF) instance.

The next-hop information is displayed for all active speaker processes in distributed mode. Each speaker displays a set of next-hops that belongs to the prefixes received by the speaker and next hops that belong to best paths that were received by other speaker processes. Use the **speaker** *speaker-id* keyword and argument to display information for only the specified speaker process.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp nexthops** command with the VRF specified:

```
RP/0/RSP0/CPU0:router# show bgp vrf all nexthops
```

```
Fri Mar 13 17:05:40.656 UTC
```

```
VRF: 900
=====
```

```
Total Nexthop Processing
  Time Spent: 0.000 secs
```

```
Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs
```

```
Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs
```

```
IPv4 Unicast is active
```

```
Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000001
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec
```

```

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I   Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201  [UR]          4294967295  0/0       1d22h (Reg)   0/3
90.0.0.2      [R][C][NL]    0           1/0       1d22h (Cri)   20/23

```

VRF: 901

=====

```

Total Nexthop Processing
  Time Spent: 0.000 secs

```

```

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

```

```

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

```

IPv4 Unicast is active

```

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000002
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

```

Nexthop Version: 1, RIB version: 1

```

Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I   Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201  [UR]          4294967295  0/0       1d22h (Reg)   0/3
91.0.0.2      [R][C][NL]    0           1/0       1d22h (Cri)   10/13

```

VRF: 902

=====

```

Total Nexthop Processing
  Time Spent: 0.000 secs

```

```

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

```

```

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

```

IPv4 Unicast is active

```

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000003
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

```

Nexthop Version: 1, RIB version: 1

show bgp nexthops

```

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I   Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201 [UR]          4294967295 0/0       1d22h (Reg)  0/3
92.0.0.2      [R][C][NL]   0           1/0       1d22h (Cri)  10/13

```

```

VRF: 903
=====

```

```

Total Nexthop Processing
  Time Spent: 0.000 secs

```

```

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

```

```

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

```

```

IPv4 Unicast is active

```

```

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000004
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

```

```

Nexthop Version: 1, RIB version: 1

```

```

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local
              I   Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201 [UR]          4294967295 0/0       1d22h (Reg)  0/3
93.0.0.2      [R][C][NL]   0           1/0       1d22h (Cri)  10/13

```

```

VRF: 904
=====

```

```

Total Nexthop Processing
  Time Spent: 0.000 secs

```

```

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

```

```

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

```

```

IPv4 Unicast is active

```

```

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000005
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

```

```

Nexthop Version: 1, RIB version: 1

```

```

Status codes: R/UR Reachable/Unreachable
              C/NC Connected/Not-connected
              L/NL Local/Non-local

```

```

I      Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201  [UR]          4294967295  0/0      1d22h (Reg)   0/3
94.0.0.2      [R][C][NL]    0           1/0      1d22h (Cri)   10/13

```

VRF: 905

=====

Total Nexthop Processing
Time Spent: 0.000 secs

Maximum Nexthop Processing
Received: 82y48w
Bestpaths Deleted: 0
Bestpaths Changed: 0
Time Spent: 0.000 secs

Last Notification Processing
Received: 1d22h
Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000006
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
C/NC Connected/Not-connected
L/NL Local/Non-local

```

I      Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201  [UR]          4294967295  0/0      1d22h (Reg)   0/3
95.0.0.2      [R][C][NL]    0           1/0      1d22h (Cri)   10/13

```

VRF: 906

=====

Total Nexthop Processing
Time Spent: 0.000 secs

Maximum Nexthop Processing
Received: 82y48w
Bestpaths Deleted: 0
Bestpaths Changed: 0
Time Spent: 0.000 secs

Last Notification Processing
Received: 1d22h
Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000007
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
C/NC Connected/Not-connected
L/NL Local/Non-local

```

I      Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
10.0.101.201  [UR]          4294967295  0/0      1d22h (Reg)   0/3

```

show bgp nexthops

```
96.0.0.2          [R][C][NL]          0          1/0          1d22h (Cri)          10/13
```

```
VRF: 907
=====
```

```
Total Nexthop Processing
  Time Spent: 0.000 secs
```

```
Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs
```

```
Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs
```

```
IPv4 Unicast is active
```

```
Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000008
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec
```

```
Nexthop Version: 1, RIB version: 1
```

```
Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I   Invalid (Policy Match Failed)
```

| Next Hop | Status | Metric | Notf | LastRIBEvent | RefCount |
|--------------|------------|------------|------|--------------|----------|
| 10.0.101.201 | [UR] | 4294967295 | 0/0 | 1d22h (Reg) | 0/3 |
| 97.0.0.2 | [R][C][NL] | 0 | 1/0 | 1d22h (Cri) | 10/13 |

```
VRF: 908
=====
```

```
Total Nexthop Processing
  Time Spent: 0.000 secs
```

```
Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs
```

```
Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs
```

```
IPv4 Unicast is active
```

```
Gateway Address Family: IPv4 Unicast
Table ID: 0xe0000009
Nexthop Count: 2
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec
```

```
Nexthop Version: 1, RIB version: 1
```

```
Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I   Invalid (Policy Match Failed)
```

| Next Hop | Status | Metric | Notf | LastRIBEvent | RefCount |
|--------------|------------|------------|------|--------------|----------|
| 10.0.101.201 | [UR] | 4294967295 | 0/0 | 1d22h (Reg) | 0/3 |
| 98.0.0.2 | [R][C][NL] | 0 | 1/0 | 1d22h (Cri) | 10/13 |

```

VRF: 909
=====

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 1d22h
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000a
Nexthop Count: 1
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1

Status codes: R/UR Reachable/Unreachable
               C/NC Connected/Not-connected
               L/NL Local/Non-local
               I Invalid (Policy Match Failed)
Next Hop      Status      Metric      Notf      LastRIBEvent RefCount
99.0.0.2      [UR]         4294967295  0/0      1d22h (Reg)  0/3
    
```

```

VRF: yellow
=====

Total Nexthop Processing
  Time Spent: 0.000 secs

Maximum Nexthop Processing
  Received: 82y48w
  Bestpaths Deleted: 0
  Bestpaths Changed: 0
  Time Spent: 0.000 secs

Last Notification Processing
  Received: 82y48w
  Time Spent: 0.000 secs

IPv4 Unicast is active

Gateway Address Family: IPv4 Unicast
Table ID: 0xe000000e
Nexthop Count: 0
Critical Trigger Delay: 0msec
Non-critical Trigger Delay: 10000msec

Nexthop Version: 1, RIB version: 1
    
```

This table describes the significant fields shown in the display.

Table 25: show bgp vrf all nexthops Field Descriptions

| Field | Description |
|-------|------------------|
| VRF | Name of the VRF. |

| Field | Description |
|---|---|
| Total Nexthop Processing Time Spent | Time spent processing trigger delays for critical and noncritical events for the VRF or address family. The time is specified in seconds. |
| Maximum Nexthop Processing | Time that has passed since the nexthop notification was received that resulted in spending the maximum amount of processing time for all notifications. |
| Last Notification Processing | Time that has passed since the last nexthop notification was received. |
| IPv4 Unicast is active. | VRF specified output that indicates the IPv4 unicast address family is active within the VRF. |
| Nexthop Count | Number of next hops for the VRF or address family. |
| Critical Trigger Delay | Configured critical trigger delay. |
| Non-critical Trigger Delay | Configured noncritical trigger delay. |
| Total Critical Notifications Received | Number of critical notifications received. |
| Total Non-critical Notifications Received | Number of noncritical notifications received. |
| Bestpaths Deleted After Last Walk | Number of best paths deleted due to the last notification. |
| Bestpaths Changed After Last Walk | Number of best paths modified due to the last notification. |
| Next Hop | IP address of the next hop. |
| Status | Status of the next hop. |
| Metric | IGP metric of the next hop. |
| Notf | Number of critical and noncritical notifications received. |
| LastRIBEvent | When the last notification was received from the RIB. |
| RefCount | The number of neighbors or prefixes that refer to the next hop in address family/all format. |
| Address Family | Name of the address family. |

Related Commands

| Command | Description |
|--|---|
| bgp redistribute-internal , on page 99 | Specifies the delay for triggering BGP next-hop calculations. |

show bgp nsr

To display Border Gateway Protocol (BGP) nonstop routing (NSR) information, use the **show bgp nsr** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6 {unicast| multicast| all|
labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vrf {vrf-name|
all}| [ipv4 {unicast| labeled-unicast}| ipv6 unicast]| vpnv6 unicast] nsr [standby]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| mdt | (Optional) Specifies multicast distribution tree (MDT) address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| vpnv6 unicast | (Optional) Specifies VPNv6 unicast address families. |
| standby | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.9.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | bgp | read |

Examples The following is sample output from the **show bgp nsr** command:

```
RP/0/RSP0/CPU0:router# show bgp nsr
Fri Jan 30 10:18:48.171 PST PDT

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs: 10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGP notified
RIB has converged: version 1
```

show bgp nsr

===== Post Failover Summary for Active instance =====

| Node | Process | Read | Write | Inbound |
|--------------|---------|--------|-------|---------|
| node0_0_CPU0 | Speaker | 146.75 | 18.90 | 3.46 |

```

Entered mode Standby Ready           : Jan 30 10:00:39
Entered mode TCP NSR Setup           : Jan 30 10:00:39
Entered mode TCP NSR Setup Done      : Jan 30 10:00:39
Entered mode TCP Initial Sync        : Jan 30 10:00:39
Entered mode TCP Initial Sync Done   : Jan 30 10:00:44
Entered mode FPBSN processing done   : Jan 30 10:00:44
Entered mode Update processing done   : Jan 30 10:00:44
Entered mode BGP Initial Sync        : Jan 30 10:00:44
Entered mode BGP Initial Sync done   : Jan 30 10:00:44
Entered mode NSR Ready                : Jan 30 10:00:44

```

Current BGP NSR state - NSR Ready achieved at: Jan 30 10:00:44
NSR State READY notified to Redcon at: Jan 30 10:16:58

NSR Post Failover Summary:

QAD Statistics:

| | | | |
|--------------------|-------|------------------------|-------|
| Messages Sent | : 512 | ACKs Received | : 512 |
| Messages Received | : 8 | ACKs Sent | : 8 |
| Send Failures | : 1 | Send ACK Failures | : 0 |
| Suspends | : 1 | Resumes | : 1 |
| Messages Processed | : 8 | Out of sequence drops: | 0 |

Postit Summary:

Total pending postit messages: 0
Neighbors with pending postits: 0

| Conv | Bestpath | TunnelUpd | Import | RIBUpd | Label | ReadWrite | LastUpd |
|------------------|----------|-----------|--------|--------|-------|-----------|---------|
| Process: Speaker | | | | | | | |
| Yes | 120 | --- | --- | 120 | 120 | 120 | 87531 |

Rib Trigger: enabled
Last RIB down event Jan 29 09:50:03.069 received
Last RIB convergence Jan 29 09:50:03.069 last ack received.

Address Family IPv4 Unicast converged in 87531 seconds

The following example shows sample output from the **show bgp nsr** command with the **standby** keyword:

RP/0/RSP0/CPU0:router# **show bgp nsr standby**

Fri Jan 30 10:18:55.654 PST PDT

```

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System: 100
Router ID: 10.1.0.1 (manually configured)
Default Cluster ID: 10.1.0.1
Active Cluster IDs: 10.1.0.1
Fast external fallover enabled
Neighbor logging is not enabled
Enforce first AS enabled
AS Path ignore is enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 180

```

```
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60
```

```
Address family: IPv4 Unicast
Dampening is not enabled
Client reflection is enabled in global config
Scan interval: 60
Main Table Version: 7034
IGP notification: IGP notified
RIB has converged: version 1
```

=====
 ===== Post Failover Summary for Standby instance =====

| Node | Process | Read | Write | Inbound |
|--------------|--------------------|------|-------|-------------------|
| node0_1_CPU0 | Speaker | 1.68 | 0.00 | 1.42 |
| Entered mode | Standby Ready | | | : Jan 30 10:00:39 |
| Entered mode | TCP Replication | | | : Jan 30 10:00:39 |
| Entered mode | TCP Init Sync Done | | | : Jan 30 10:00:44 |
| Entered mode | NSR Ready | | | : Jan 30 10:00:44 |

QAD Statistics:

| | | | |
|--------------------|-------|--------------------|-----------------|
| Messages Sent | : 9 | ACKs Received | : 9 |
| Messages Received | : 512 | ACKs Sent | : 512 |
| Send Failures | : 0 | Send ACK Failures | : 0 |
| Suspends | : 0 | Resumes | : 0 |
| Messages Processed | : 512 | Standby init drops | : 0 |
| drops: 0 | | | Out of sequence |

Postit Summary:

```
Total pending postit messages: 0
Neighbors with pending postits: 0
```

| Conv | Bestpath | TunnelUpd | Import | RIBUpd | Label | ReadWrite | LastUpd |
|------------------|------------|-----------|--------|------------|------------|------------|---------|
| Process: Speaker | | | | | | | |
| Yes | 1233338444 | --- | --- | 1233338444 | 1233338444 | 1233338444 | --- |

```
Rib Trigger: enabled
Last RIB down event Jan 29 09:50:17.308 received
Last RIB convergence Jan 29 09:50:17.308 last ack received.
```

Related Commands

| Command | Description |
|--|--|
| nsr (BGP), on page 210 | Activates Border Gateway Protocol (BGP) nonstop routing (NSR). |

show bgp paths

To display all the Border Gateway Protocol (BGP) paths in the database, use the **show bgp paths** command in EXEC mode.

show bgp paths [**detail**] [**debug**] [**regexp** *regular-expression*]

Syntax Description

| | |
|---|---|
| detail | (Optional) Displays detailed attribute information. |
| debug | (Optional) Displays attribute process ID, hash bucket, and hash chain ID attribute information. |
| regexp <i>regular-expression</i> | (Optional) Specifies an autonomous system path that matches the regular expression. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp paths** command to display information about AS paths and the associated attributes with which the paths were received.

If no options are specified, all stored AS paths are displayed with the number of routes using each path.



Note

The AS path information is stored independently of the address family, making it possible that routes from different address families could be using the same path.

Use the *regular-expression* argument to limit the output to only those paths that match the specified regular expression. See the *Cisco ASR 9000 Series Aggregation Services Router Getting Started Guide* for information on regular expressions.

Use the **detail** keyword to display detailed information on the attributes stored with the AS path.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp paths** command:

```
RP/0/RSP0/CPU0:router# show bgp paths detail

Proc  Attributes                               Refcount    Metric Path
Spk 0  ORG AS LOCAL                             7           0 i
Spk 0  ORG AS LOCAL COMM EXTCOMM              3           0 21 i
Spk 0  MET ORG AS                             3           55 2 i
Spk 0  ORG AS                                 3           0 2 10 11 i
Spk 0  ORG AS COMM                            3           0 2 10 11 i
Spk 0  MET ORG AS ATOM                        3           2 2 3 4 ?
Spk 0  MET ORG AS                             3           1 2 3 4 e
Spk 0  MET ORG AS                             3           0 2 3 4 i
```

This table describes the significant fields shown in the display.

Table 26: show bgp paths Field Descriptions

| Field | Description |
|------------|---|
| Proc | ID of the process in which the path is stored. This is always "Spk 0." |
| Attributes | Attributes that are present. The following may appear: MET—Multi Exit Discriminator (MED) attribute is present. ORG—Origin attribute is present. AS—AS path attribute is present. LOCAL—Local preference attribute is present. AGG—Aggregator attribute is present. COMM—Communities attribute is present. ATOM—Atomic aggregate attribute is present. EXTCOMM—Extended communities attribute is present. |

| Field | Description |
|------------|---|
| NeighborAS | <p>Autonomous system number of the neighbor, or 0, if the path information originated locally.</p> <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Refcount | Number of routes using a path. |
| Metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| Path | <p>Autonomous system path to the destination network. At the end of the path is the origin code for the path:</p> <p>i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command.</p> <p>e—Path originated from an Exterior Gateway Protocol (EGP).</p> <p>?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.</p> |

show bgp policy

To display information about Border Gateway Protocol (BGP) advertisements under a proposed policy, use the **show bgp policy** command in EXEC mode.

show bgp policy

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd <i>rd-address</i> | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| neighbor | (Optional) Previews advertisements for a single neighbor. |
| <i>ip-address</i> | (Optional) IP address of a single neighbor. |
| sent-advertisements | (Optional) Displays the routes that have been advertised to neighbors. If a route has not yet been advertised to the neighbor, it is not shown. |
| route-policy | (Optional) Displays advertisements for an output route policy. |
| <i>route-policy-name</i> | (Optional) Name of the route policy. |

| | |
|----------------|--|
| summary | (Optional) Displays a summary of the BGP advertisements. |
|----------------|--|

Command Default

Advertisements for all neighbors are displayed if the **neighbor** *ip-address* keyword and argument are not specified. If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp policy** command to display routes that would be advertised to neighbors under a proposed policy. Unlike in the **show bgp advertised** command, the information displayed reflects any modifications made to the routes when executing the specified policy.

Use the **neighbor** keyword to limit the output to routes advertised to a particular neighbor. Use the **sent-advertisements** keyword to change the output in two ways:

- If a policy is not specified explicitly, any policy configured on the neighbor (using the **route-policy (BGP)** command) is executed before displaying the routes.
- Only routes that have already been advertised to the neighbor (and not withdrawn) are displayed. Routes that have not yet been advertised are not displayed.

Use the **summary** keyword to display abbreviated output.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp policy** command with the **summary** keyword in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp policy summary
Network      Next Hop      From           Advertised to
172.16.1.0/24 10.0.101.1    10.0.101.1    10.0.101.2
                                     10.0.101.3

172.17.0.0/16 0.0.0.0       Local          10.0.101.1
                                     10.0.101.2
                                     10.0.101.3
```

This table describes the significant fields shown in the display.

Table 27: show bgp policy summary Field Descriptions

| Field | Description |
|-----------------|--|
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| From | IP address of the peer that advertised this route. |
| Local | Indicates the route originated on the local system. |
| Local Aggregate | Indicates the route is an aggregate created on the local system. |
| Advertised to | Indicates the neighbors to which this route was advertised. |

The following is sample output from the **show bgp policy** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp policy
11.0.0.0/24 is advertised to 10.4.101.1
  Path info:
    neighbor: Local           neighbor router id: 10.4.0.1
    valid local best
  Attributes after inbound policy was applied:
    next hop: 0.0.0.0
    MET ORG AS
```

```

    origin: IGP metric: 0
    aspath:
Attributes after outbound policy was applied:
  next hop: 10.4.0.1
  MET ORG AS
  origin: IGP metric: 0
  aspath: 1

11.0.0.0/24 is advertised to 10.4.101.2
Path info:
  neighbor: Local          neighbor router id: 10.4.0.1
  valid local best
Attributes after inbound policy was applied:
  next hop: 0.0.0.0
  MET ORG AS
  origin: IGP metric: 0
  aspath:
Attributes after outbound policy was applied:
  next hop: 10.4.0.1
  MET ORG AS
  origin: IGP metric: 0
  aspath:

11.0.0.0/24 is advertised to 10.4.101.3
Path info:
  neighbor: Local          neighbor router id: 10.4.0.1
  valid local best
Attributes after inbound policy was applied:
  next hop: 0.0.0.0
  MET ORG AS
  origin: IGP metric: 0
  aspath:
Attributes after outbound policy was applied:
  next hop: 10.4.0.1
  MET ORG AS
  origin: IGP metric: 0
  aspath:

12.0.0.0/24 is advertised to 10.4.101.2
Path info:
  neighbor: 10.4.101.1     neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

12.0.0.0/24 is advertised to 10.4.101.3
Path info:
  neighbor: 10.4.101.1     neighbor router id: 10.4.101.1
  valid external best
Attributes after inbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath: 2 3 4
Attributes after outbound policy was applied:
  next hop: 10.4.101.1
  ORG AS
  origin: IGP neighbor as: 2
  aspath:2 3 4

```

This table describes the significant fields shown in the display.

Table 28: show bgp policy Field Descriptions

| Field | Description |
|-------------------------------------|--|
| Is advertised to | IP address of the peer to which this route is advertised. If the route is advertised to multiple peers, information is shown separately for each peer. |
| neighbor | IP address of the peer that advertised this route, or one of the following: Local—Route originated on the local system. Local Aggregate—Route is an aggregate created on the local system. |
| neighbor router id | BGP identifier for the peer, or the local system if the route originated on the local system. |
| Not advertised to any peer | Indicates the no-advertise well-known community is associated with this route. Routes with this community are not advertised to any BGP peers. |
| Not advertised to any EBGp peer | Indicates the no-export well-known community is associated with this route. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router. |
| Not advertised outside the local AS | Indicates the local-AS well-known community is associated with this route. Routes with this community value are not advertised outside the local autonomous system or confederation boundary. |
| (Received from a RR-client) | Path was received from a route reflector client. |
| (received-only) | Path is not used for routing purposes. It is used to support soft reconfiguration, and records the path attributes before inbound policy was applied to a path received from a peer. A path marked “received-only” indicates that either the path was dropped by inbound policy, or that a copy of path information was created and then modified for routing use. |
| (received & used) | Indicates that the path is used both for soft reconfiguration and routing purposes. A path marked “(received & used)”, implies the path information was not modified by inbound policy. |
| valid | Path is valid. |
| redistributed | Path is locally sourced through redistribution. |

| Field | Description |
|---|--|
| aggregated | Path is locally sourced through aggregation. |
| local | Path is locally sourced through the network command. |
| confed | Path was received from a confederation peer. |
| best | Path is selected as best. |
| multipath | Path is one of multiple paths selected for load-sharing purposes. |
| dampinfo | Indicates dampening information: Penalty—Current penalty for this path. Flapped—Number of times the route has flapped. In—Time (hours:minutes:seconds) since the network first flapped. Reuse in—Time (hours:minutes:seconds) after which the path is available. This field is displayed only if the path is currently suppressed. |
| Attributes after inbound policy was applied | Displays attributes associated with the received route, after any inbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present. |

| Field | Description |
|-------------|---|
| origin | <p>Origin of the path:</p> <p>IGP—Path originated from an Interior Gateway Protocol (IGP) and was sourced by BGP using a network or aggregate-address command.</p> <p>EGP—Path originated from an Exterior Gateway Protocol.</p> <p>incomplete—Origin of the path is not clear; in example, a route that is redistributed into BGP from an IGP.</p> |
| neighbor as | First autonomous system (AS) number in the AS path. |
| aggregator | Indicates that the path was received with the aggregator attribute. The AS number and router-id of the system that performed the aggregation are shown. |
| metric | Value of the interautonomous system metric, otherwise known as the MED metric. |
| localpref | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system |
| aspath | AS path associated with the route. |
| community | <p>Community attributes associated with the path. Community values are displayed in AA:NN format, except for the following well-known communities:</p> <p>Local-AS—Community with value 4294967043 or hex 0xFFFFFFFF03. Routes with this community value are not advertised outside the local autonomous system or confederation boundary.</p> <p>no-advertise—Community with value 4294967042 or hex 0xFFFFFFFF02. Routes with this community value are not advertised to any BGP peers.</p> <p>no-export—Community with value 4294967041 or hex 0xFFFFFFFF01. Routes with this community are not advertised to external BGP peers, even if those peers are in the same confederation as the local router.</p> |

| Field | Description |
|--|---|
| Extended community | Extended community attributes associated with the path. For known extended community types, the following codes may be displayed: RT—Route target community SoO—Site of Origin community LB—Link Bandwidth community |
| Originator | Router ID of the originating router when route reflection is used. |
| Cluster lists | Router ID or cluster ID of all route reflectors through which the route has passed. |
| Attributes after outbound policy was applied | Displays attributes associated with the received route, after any outbound policy has been applied. AGG—Aggregator attribute is present. AS—AS path attribute is present. ATOM—Atomic aggregate attribute is present. COMM—Communities attribute is present. EXTCOMM—Extended communities attribute is present. LOCAL—Local preference attribute is present. MET—Multi Exit Discriminator (MED) attribute is present. next hop—IP address of the next system used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. ORG—Origin attribute is present. |

Related Commands

| Command | Description |
|--|---|
| route-policy (BGP) , on page 249 | Applies an inbound or outbound routing policy to a neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

| Command | Description |
|--|--|
| show bgp advertised, on page 293 | Displays routes advertised to neighbors. |
| show bgp neighbors, on page 353 | Displays information about the TCP and BGP connections to neighbors. |
| show bgp route-policy, on page 431 | Displays BGP information about networks that match an outbound route policy. |

show bgp process

To display Border Gateway Protocol (BGP) process information, use the **show bgp process** command in EXEC mode.

```
show bgp [ipv4| {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6| {unicast| multicast| all|
labeled-unicast}| all| {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vpnv6 unicast]
process [performance-statistics] [detail]
```

Syntax Description

| | |
|--------------------------------|--|
| ipv4 | (Optional) Specifies IP Version 4. |
| unicast | (Optional) Specifies the unicast subaddress family. |
| multicast | (Optional) Specifies the multicast subaddress family. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| performance- statistics | (Optional) Displays performance statistics relative to the work done by the specified process. |
| detail | (Optional) Specifies detailed process information. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 3.9.0 | Asplain format for 4-byte autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.0 | The command output was modified to include information from BGP additional paths send and receive capability configurations. |
| Release 4.3.0 | The command output was modified to include information from update wait-install configuration. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp process** command to display status and summary information for the Border Gateway Protocol (BGP) process. The output shows various global and address family-specific BGP configurations. A summary of the number of neighbors, update messages, and notification messages sent and received by the process is also displayed.

Use the **detail** keyword to display detailed process information. The detailed process information shows the memory used by each of various internal structure types.

Use the **performance-statistics** keyword to display a summary or detail of work done by the BGP processes. The summary display shows the real time spent performing certain operations and the time stamps for state transitions during initial convergence.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp process** command:

```
RP/0/RSP0/CPU0:router# show bgp process

BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
```

show bgp process

```

Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Scan interval: 60
Main Table Version: 150
IGP notification: IGP notified

Node      Process      Nbrs Estab Rst Upd-Rcvd Upd-Sent Nfn-Rcvd Nfn-Sent
node0_0_CPU0 Speaker      3     2   1     20      10       0        0

```

This table describes the significant fields shown in the display.

Table 29: show bgp process Field Descriptions

| Field | Description |
|---------------------|--|
| BGP is operating in | Indicates BGP is operating in standalone mode. This is the only supported mode. |
| Autonomous System | Autonomous system number for the local system. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If no global ID is available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. This is controlled by the default-metric command. |

| Field | Description |
|--|---|
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. This is controlled by the bgp fast-external-fallover disable command. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the bgp log neighbor changes disable command. |
| Enforce first AS enabled | Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. This is controlled by the bgp enforce-first-as disable command. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing Multi Exit Discriminator (MED) metric values are treated as worst in the route selection algorithm. This is controlled by the bgp bestpath med missing-as-worst command. |
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. This is controlled by the bgp default local-preference command. |
| Default keepalive | Default keepalive interval. This is controlled by the timers bgp command. |

| Field | Description |
|--------------------------|--|
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| Dampening | Indicates whether dampening is enabled for the specified address family. This is controlled by the dampening command. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command in address family configuration mode. |
| Main Table Version | Last version of the BGP database that was installed into the main routing table. |
| IGP notification | Indicates whether Interior Gateway Protocols (IGP) have been notified of BGP convergence for the specified address family. |
| Node | Node on which the process is executing. |
| Process | Type of BGP process. |
| Speaker | Speaker process. A speaker process is responsible for receiving, processing, and sending BGP messages to configured neighbors. |

| Field | Description |
|----------|--|
| Nbrs | Number of neighbors for which the process is responsible. |
| Estab | Number of neighbors that have connections in the established state for this process. |
| Rst | Number of times this process was restarted. |
| Upd-Rcvd | Number of update messages received by the process. |
| Upd-Sent | Number of update messages sent by the process. |
| Nfn-Rcvd | Number of notification messages received by the process. |
| Nfn-Sent | Number of notification messages sent by the process. |

The following is sample output from the **show bgp process** command with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show bgp all all process detail
```

```
BGP Process Information
BGP is operating in STANDALONE mode
Autonomous System: 1
Router ID: 10.0.0.5 (manually configured)
Cluster ID: 10.0.0.5
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
Default local preference: 100
Default keepalive: 60
Update delay: 120
Generic scan interval: 60

BGP Speaker process: 0, location node0_0_0
Neighbors: 3, established: 2

                Sent           Received
Updates:         3             15
Notifications:  0             0

                Number         Memory Used
Attributes:     12            1104
AS Paths:       10            400
Communities:    2            1080
Extended communities: 1         40
Route Reflector Entries: 0         0
Route-map Cache Entries: 0         0
Filter-list Cache Entries: 0         0
Next Hop Cache Entries: 2         80
Update messages queued: 0

Address family: IPv4 Unicast
Dampening is enabled
Client reflection is enabled
Main Table Version: 12
IGP notification: IGP notified

State: normal mode.
BGP Table Version: 12
```

show bgp process

```
Network Entries: 15, Soft Reconfig Entries: 0
Dampened Paths: 0, History Paths: 9
```

```

                Allocated      Freed
Prefixes:      15              0
Paths:         19              0

                Number        Memory Used
Prefixes:      15             1230
Paths:         19             760

```

This table describes the significant fields shown in the display.

Table 30: show bgp process detail Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP is operating in | Indicates whether BGP is operating in standalone mode. |
| Autonomous System | Autonomous system number for the local system. |
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | Cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. |
| Enforce first AS enabled | Indicates that strict checking of the first autonomous system (AS) number in paths received from external BGP peers is enabled. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled by the bgp bestpath med missing-as-worst command. |

| Field | Description |
|--|--|
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. |
| Default keepalive | Default keepalive interval. This is controlled by the timers bgp command. |
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: <ul style="list-style-type: none"> • bgp graceful-restart • bgp graceful-restart purge-time • bgp graceful-restart stalepath-time • bgp graceful-restart restart-time • bgp graceful-restart graceful-reset |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This is controlled by the bgp scan-time command. |
| BGP Speaker Process | Speaker process responsible for receiving, processing and sending BGP messages. |

| Field | Description |
|-------------------------|--|
| Node | Node on which the specified process is executing. |
| Neighbors | Number of neighbors for which the specified process is responsible. |
| established | Number of neighbors that have connections in the established state for the specified process. |
| Updates | Number of update messages sent and received by the specified process. |
| Notifications | Number of notification messages sent and received by the specified process. |
| Attributes | Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information. |
| AS Paths | Number of unique autonomous system paths stored in the specified process and the amount of memory used by the AS path information. |
| Communities | Number of unique sets of community information stored in the specified process and the amount of memory used by them. |
| Extended communities | Number of unique sets of extended community information stored in the specified process and the amount of memory used by them. |
| Route Reflector Entries | Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them. |
| Nexthop Entries | Number of entries and memory usage for cached next-hop information. |
| Update messages queued | Total number of update messages queued to be sent across all neighbors for which the specified process is responsible. |
| Address family | Specified address family. |
| Dampening | Indicates whether dampening is enabled for the specified address family. |

| Field | Description |
|--------------------|---|
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command. |
| Main Table Version | Last version of the local BGP database for the specified address family that was injected into the main routing table. |
| IGP notification | Indicates whether IGP have been notified of BGP convergence for the specified address family. |
| RIB has converged | Indicates whether the main routing table version has converged and the version at which it converged. |
| State | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.</p> |
| BGP Table Version | Last version used in the BGP database for received routes. |

| Field | Description |
|-------------------------------|--|
| Attribute download | Indicates whether the RIB attribute download is enabled. |
| Network Entries | Number of sets of prefix information held in the specified BGP process for the specified address family. |
| Soft Reconfig Entries | Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration. |
| Dampened Paths | Number of routes that are suppressed due to dampening for the specified address family. |
| History Paths | Number of routes that are currently withdrawn, but are being maintained to preserve dampening information. |
| Prefixes (Allocated/Freed) | Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Paths (Allocated/Freed) | Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Prefixes (Number/Memory Used) | Number of sets of prefix information currently allocated for the specified address family, and the amount of memory used by them. |
| Paths (Number/Memory Used) | Number of sets of route information currently allocated for the specified address family, and the amount of memory used by them. |

The following is sample output from the **show bgp process** command with the **performance-statistics** keyword:

```
RP/0/RSP0/CPU0:router# show bgp process performance-statistics detail
```

```
BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2
```

```

                Sent           Received
Updates:         20             20
Notifications:  0              0
```

```

                Number         Memory Used
Attributes:     2              184
AS Paths:       2              48
Communities:    0              0
Extended communities: 0          0
```

```

Route Reflector Entries: 0 0
Route-map Cache Entries: 0 0
Filter-list Cache Entries: 0 0
Next Hop Cache Entries: 2 80
Update messages queued: 0 0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)

State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0

Prefixes: Allocated Freed
Paths: 149 0
        200 0

Prefixes: Number Memory Used
Paths: 149 12516
        200 8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

Update wait-install enabled:
ack request 2, ack rcvd 2, slow ack 0
Max time for batched RIB update:
update time 0 secs, 1752000 nsecs
install time 0 secs, 36391000 nsecs
routes installed 6, modified 0, withdrawn 2
start version 1, target version 15

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
    
```

This table describes the significant fields shown in the display.

Table 31: show bgp process performance-statistics Field Descriptions

| Field | Description |
|---------------------|--|
| BGP is operating in | Indicates whether BGP is operating in standalone mode. |
| Autonomous system | Autonomous system number for the local system. |

| Field | Description |
|--|--|
| Router ID | BGP identifier assigned to the local system. If this is explicitly configured using the bgp router-id command, “manually configured” is displayed. If the router ID is not manually configured, it is determined from a global router ID. If the global ID is not available, the router ID is shown as 0.0.0.0. |
| Confederation ID | Confederation identifier for the local system. |
| Cluster ID | The cluster identifier for the local system. If this is manually configured using the bgp cluster-id command, “manually configured” is displayed. |
| Default metric | Default metric. |
| Fast external fallover enabled | Indicates whether fast external fallover is enabled. |
| Neighbor logging enabled | Indicates whether logging of peer connection up and down transitions is enabled. This is controlled by the bgp log neighbor changes disable command. |
| Enforce first AS enabled | Indicates that strict checking of the first AS number in paths received from external BGP peers is enabled. |
| iBGP to IGP redistribution | Indicates internal redistribution is enabled using the bgp redistribution-internal command. |
| Treating missing MED as worst | Indicates missing MED metric values are treated as worst in the route selection algorithm. This is controlled using the bgp bestpath med missing-as-worst command. |
| Always compare MED is enabled | Indicates that the MED is always used during the route selection algorithm, even when paths are received from external BGP neighbors in different autonomous systems. This setting is controlled by the bgp bestpath med always command. |
| AS Path ignore is enabled | Indicates that the AS path length is ignored by the route selection algorithm. This is controlled by the bgp bestpath as-path ignore command. |
| Comparing MED from confederation peers | Indicates that the MED values are used in the route selection algorithm when comparing routes received from confederation peers. This is controlled by the bgp bestpath med confed command. |

| Field | Description |
|------------------------------------|--|
| Comparing router ID for eBGP paths | Indicates that the router ID is used as a tiebreaker by the route selection algorithm when comparing identical routes received from different external BGP neighbors. This is controlled by the bgp bestpath compare-routerid command. |
| Default local preference | Default local preference value used for BGP routes. |
| Default keepalive | Default keepalive interval. This setting is controlled by the timers bgp command. |
| Graceful restart enabled | Indicates that the graceful restart capability is enabled. The configuration commands affecting graceful restart behavior are: bgp graceful-restart , bgp graceful-restart purge-time , bgp graceful-restart stalepath-time , bgp graceful-restart restart-time , and bgp graceful-restart graceful-reset . |
| Update delay | Maximum time that a BGP process stays in read-only mode. |
| Generic scan interval | Interval (in seconds) between BGP scans for address family-independent tasks. This setting is controlled by the bgp scan-time command in router configuration mode. |
| Address family | Specified address family. |
| Dampening | Indicates whether dampening is enabled for the specified address family. |
| Client reflection | Indicates whether client-to-client route reflection is enabled for the specified address family. This is controlled by the bgp client-to-client reflection disable command. |
| Scan interval | Interval (in seconds) between BGP scans for the given address family. This is controlled by the bgp scan-time command. |
| Main Table Version | Last version of the local BGP database for the specified address family that was injected into the main routing table. |
| IGP notification | Indicates whether IGP have been notified of BGP convergence for the specified address family. |
| Node | Node on which the process is executing. |

| Field | Description |
|---|--|
| Process | BGP process. |
| Speaker | Speaker process. The speaker process is responsible for receiving, processing and sending BGP messages. |
| Read | Real time (in seconds) spent reading messages from peers by this process. |
| Write | Real time (in seconds) spent writing messages to peers by this process. |
| Inbound | The real time (in seconds) spent processing messages read from peers by this process. |
| Config | Real time (in seconds) spent processing configuration commands by this process. |
| Data | Real time (in seconds) spent providing operational data by this process. |
| Conv | Indicates whether the process has converged after the initial update. |
| Nbr Estab | Time stamp (in seconds) recording the time when the first neighbor became established. |
| Bestpath | Time stamp (in seconds) recording the time the best-path calculation mode was entered. |
| RIB Inst | Time stamp (in seconds) recording the time RIB update mode was entered. |
| Read/Write | Time stamp (in seconds) recording the time normal mode was entered. |
| Last Upd | Time stamp (in seconds) recording the time the last update was sent to a neighbor. |
| Address Family IPv4 Unicast converged in <i>n</i> seconds | Indicates that BGP has reached initial convergence for the IPv4 unicast address family. The time taken for convergence is shown. |
| Address Family IPv6 Multicast converged in <i>n</i> seconds | Indicates that BGP has reached initial convergence for the IPv6 multicast address family. The time taken for convergence is shown. |
| Update wait-install enabled | Indicates the update wait-install was configured. |

The following is sample output from the **show bgp process** command with the **performance-statistics** and **detail** keywords:

```
RP/0/RSP0/CPU0:router# show bgp process performance-statistics detail

BGP Speaker process: 0, Node: node0_0_CPU0
Restart count: 2
Neighbors: 3, established: 2

Updates:                Sent           Received
Notifications:         0             0

Attributes:             Number       Memory Used
AS Paths:               2           184
Communities:            0           48
Extended communities:  0           0
Route Reflector Entries: 0           0
Route-map Cache Entries: 0           0
Filter-list Cache Entries: 0           0
Next Hop Cache Entries: 2           80
Update messages queued: 0

Read 14 messages (1142 bytes) in 12 calls (time spent: 0.024 secs)
Read throttled 0 times
Processed 14 inbound messages (time spent: 0.132 secs)
Wrote 2186 bytes in 24 calls (time spent: 0.024 secs)
Processing write list: wrote 18 messages in 4 calls (time spent: 0.000 secs)
Processing write queue: wrote 10 messages in 20 calls (time spent: 0.000 secs)
Socket setup (LPTS): 4 calls (time spent: 0.010 secs)
Configuration: 1 requests (time spent: 0.002 secs)
Operational data: 9 requests (time spent: 0.026 secs)

State: normal mode.
BGP Table Version: 150
Network Entries: 149, Soft Reconfig Entries: 0

Prefixes:               Allocated   Freed
Paths:                  149         0
                       200         0

Prefixes:               Number       Memory Used
Paths:                  149         12516
                       200         8000

Updates generated: 149 prefixes in 8 messages from 2 calls (time spent: 0.046 secs)
Scanner: 2 scanner runs (time spent: 0.008 secs)
RIB update: 1 rib update runs, 149 prefixes installed (time spent: 0.024 secs)
Process has converged for IPv4 Unicast.

First neighbor established: 1082604050s
Entered DO_BESTPATH mode: 1082604055s
Entered DO_RIBUPD mode: 1082604055s
Entered Normal mode: 1082604055s
Latest UPDATE sent: 1082604056s
```

This table describes the significant fields shown in the display.

Table 32: show bgp process performance-statistics detail Field Descriptions

| Field | Description |
|---------|------------------------|
| Process | The specified process. |

| Field | Description |
|---------------------------|--|
| Location | Node in which the specified process is executing. |
| Neighbors | Number of neighbors for which the specified process is responsible. |
| established | Number of neighbors that have connections in the established state for the specified process. |
| Updates | Number of update messages sent and received by the specified process. |
| Notifications | Number of notification messages sent and received by the specified process. |
| Attributes | Number of unique sets of attribute information stored in the specified process and the amount of memory used by the attribute information. |
| AS Paths | Number of unique autonomous system paths stored in the specified process, and the amount of memory used by the AS path information. |
| Communities | Number of unique sets of community information stored in the specified process and the amount of memory used by them. |
| Extended communities | Number of unique sets of extended community information stored in the specified process and the amount of memory used by them. |
| Route Reflector Entries | Number of unique sets of route reflector information stored in the specified process and the amount of memory used by them. |
| Route-map Cache Entries | Number of entries and memory usage for cached results for applying a route map. |
| Filter-list Cache Entries | Number of entries and memory usage for cached results for applying an AS path filter list. |
| Next Hop Cache Entries | Number of entries and memory usage for cached next-hop information. |
| Update messages queued | Number of update messages queued to be sent across all neighbors for which the specified process is responsible. |

| Field | Description |
|------------------------|---|
| Read | Indicates the number of messages read by the process, the total size of read messages, the number of read operations performed, and the real time spent by the process performing read operations. |
| Read throttled | Number of times that reading from TCP has been throttled due to a backlog of messages read but not processed. |
| inbound messages | Number of read messages that have been processed and the real time spent processing inbound messages. |
| Wrote | Amount of data that has been written by the process, the number of write operations performed, and the real time spent by the process performing write operations. |
| Processing write list | Number of messages written from write lists, the number of times the write list has been processed, and the real time spent processing the write list. Note Write lists typically contain only update messages. |
| Processing write queue | Number of messages written from write queues, number of times the write queue has been processed, and the real time spent processing the write queue. |
| Socket setup | Number of socket setup operations performed and the real time spent during socket setup operations. |
| Configuration | Number of configuration requests received by the process and the real time spent processing configuration requests. |
| Operational data | Number of requests for operational data (for show commands) received by the process and the real time spent processing operation data requests |

| Field | Description |
|----------------------------|---|
| State | <p>BGP system state for the specified address family and process. This may be one of the following:</p> <p>read-only mode—Initial set of updates is being recovered. In this mode, route selection is not performed, routes are not installed in the global RIB, and updates are not advertised to peers.</p> <p>best-path calculation mode—Route selection is being performed for the routes that were received while in read-only mode.</p> <p>import mode—Routes are imported from one VRF to another VRF once the best paths are calculated. This mode is supported in VPNv4 unicast address family mode.</p> <p>RIB update mode—Routes that were selected in best-path calculation mode are being installed in the global RIB.</p> <p>label allocation mode: Labels are allocated for the received prefixes based on the requirement.</p> <p>normal mode—Best paths are sent to the peers for routes that exist in the RIB. The route selection, import processing, RIB updates, and label allocation are performed as new updates are received.</p> |
| BGP Table Version | Last version used in the BGP database for received routes. |
| Network Entries | Number of sets of prefix information held in the specified BGP process for the specified address family. |
| Soft Reconfig Entries | Number of sets of prefix information that are present only for the purpose of supporting soft reconfiguration. |
| Dampened Paths | Number of routes that are suppressed due to dampening for the specified address family. |
| History Paths | Number of routes that are currently withdrawn, but are being maintained to preserve dampening information. |
| Prefixes (Allocated/Freed) | Number of sets of prefix information for the specified address family that have been allocated and freed during the lifetime of the process. |

| Field | Description |
|-------------------------------|--|
| Paths (Allocated/Freed) | Number of sets of route information for the specified address family that have been allocated and freed during the lifetime of the process. |
| Prefixes (Number/Memory Used) | Number of sets of prefix information currently allocated for the specified address family and amount of memory used by them. |
| Paths (Number/Memory Used) | Number of sets of route information currently allocated for the specified address family and amount of memory used by them. |
| Updates generated | Number of prefixes for which updates have been generated, the number of messages used to advertise the updates, the number of update generation runs performed, and the real time spent generating updates for the specified address family. |
| Scanner | Number of times the scanner has run for the specified address family and real time spent in scanner processing. |
| RIB Update | Number of global routing information base update runs performed for the specified address family, number of prefixes installed, withdrawn, or modified in the global RIB during these runs, and real time spent performing these runs. |
| Process has converged | Indicates whether the process has reached initial convergence for the specified address family. |
| First neighbor established | Time stamp (in seconds) recording the time the first neighbor in the process was established. |
| Entered DO_BESTPATH mode | Time stamp (in seconds) recording the time best-path calculation mode was entered. |
| Entered DO_RIBUPD mode | Time stamp (in seconds) recording the time RIB update mode was entered. |
| Entered Normal mode | Time stamp (in seconds) recording the time normal mode was entered. |
| Last UPDATE sent | Time stamp (in seconds) recording the time the last update was sent to a neighbor. |

The following is sample output from the **show bgp vpnv4 unicast process performance-statistics detail** command:

```
RP/0/RSP0/CPU0:router# show bgp vpnv4 unicast process performance-statistics detail
BGP Speaker process: 0, Node: node0_8_CPU0 Restart count: 1
      Total          Nbrs Estab/Cfg
Default VRFs:          1             4/12
Non-Default VRFs:     1009          1082/1337

      Sent          Received
Updates:              362259         5688505
Notifications:        14             0

      Number        Memory Used
Attributes:           14896         2979200
AS Paths:             17             1100
Communities:          3             120
Extended communities: 1849         124440
Route Reflector Entries: 417         25020
Nexthop Entries:     2941         539572
Update messages queued: 0

      Alloc          Free
Pool 210:             28955629        28955628
Pool 310:             363103          363103
Pool 600:             4931162         4931162
Pool 1100:            104693          104693
Pool 4300:            799374          799374

Read 34755745 messages (3542094326 bytes) in 30528983 calls (time spent: 6427.769 secs)
Read partly throttled 1506 times
  Read 14 times after crossing lower threshold Processed 5836892 inbound update messages
  (time spent: 6229.512 secs)
Wrote 825719955 bytes in 29272669 calls (time spent: 2318.472 secs)
Processing sub-group: wrote 861402 messages in 1113810 calls (time spent: 145.446 secs)
Processing write queue: wrote 6288 messages in 20498 calls (time spent: 0.039 secs)
Socket setup (LPTS): 0 calls (time spent: 0.000 secs)
event_file_attach calls: Input 8769, Output 2810, Input-output 0
Configuration: 989 requests (time spent: 0.046 secs) Operational data: 92396 requests (time
spent: 98.864 secs)
Current Clock Time: not set Update Generation master timer:
  id: 0, time left: 0.0 sec, last processed: not set
  expiry time of parent node: not set
IO master timer:
  id: 0, time left: 0.0 sec, last processed: not set
  expiry time of parent node: not set

Address Family: VPNv4 Unicast
State: Normal mode.
BGP Table Version: 23211188
Attribute download: Disabled
Soft Reconfig Entries: 0

      Last 8 Triggers          Ver          Tbl Ver
Label Thread
Jun 18 05:31:39.120 23211188 23211188
Jun 18 05:31:35.274 23211188 23211188
Jun 18 05:31:34.340 23211187 23211188
Jun 18 05:31:34.189 23211186 23211187
Jun 18 05:31:29.120 23211186 23211186
Jun 18 05:31:28.861 23211186 23211186
Jun 18 05:31:19.640 23211186 23211186
Jun 18 05:31:19.272 23211186 23211186
Total triggers: 639526

Import Thread
Jun 18 05:31:39.120 23211188 23211188
Jun 18 05:31:35.274 23211188 23211188
Jun 18 05:31:34.340 23211187 23211188
Jun 18 05:31:34.189 23211186 23211187
Jun 18 05:31:29.120 23211186 23211186
Jun 18 05:31:28.861 23211186 23211186
```

```

Jun 18 05:31:19.640 23211186 23211186
Jun 18 05:31:19.272 23211186 23211186
Total triggers: 689177

RIB Thread Jun 18 05:31:39.146 23211188 23211188
Jun 18 05:31:35.299 23211188 23211188
Jun 18 05:31:34.525 23211187 23211188
Jun 18 05:31:34.494 23211186 23211188
Jun 18 05:31:34.340 23211186 23211188
Jun 18 05:31:34.255 23211186 23211188
Jun 18 05:31:29.146 23211186 23211186
Jun 18 05:31:28.886 23211186 23211186
Total triggers: 668084

Update Thread Jun 18 05:31:39.171 --- 23211188
Jun 18 05:31:35.324 --- 23211188
Jun 18 05:31:34.558 --- 23211188
Jun 18 05:31:34.521 --- 23211188
Jun 18 05:31:34.327 --- 23211188
Jun 18 05:31:29.170 --- 23211186
Jun 18 05:31:28.910 --- 23211186
Jun 18 05:31:19.690 --- 23211186
Total triggers: 660143

Remote Prefixes: Allocated Freed
Remote Paths: 3150972 2885064
7639074 7118286

Local Prefixes: 3760870 3425614
Local Paths: 7892100 7595657

Remote Prefixes: Number Mem Used
Remote Paths: 265908 29781696
Remote RDs: 520788 24997824
12424 2832672

Local Prefixes: 335256 37548672
Local Paths: 296443 14229264
Local RDs: 1009 230052

Total Prefixes: 601164 67330368
Total Paths: 817231 39227088
Imported Paths: 265675 12752400
Total RDs: 13433 3062724
Same RDs: 0 0

Update Groups: 3 Subgroups: 2
Updates generated: 1438448 prefixes in 67375 messages from 181564 calls (time spent: 6779.576
secs)
Scanner: 0 scanner runs (time spent: 0.000 secs) RIB update: 0 rib update runs, 0 prefixes
installed, 0 modified,
0 prefixes removed (time spent: 0.000 secs) RIB table update: 0 table deletes,
0 table invalid, 3526736604 table skip,
0 no local label, 0 rib retries Process has not converged for VPNv4 Unicast.

First neighbor established: Jun 11 08:32:10
Entered DO_BESTPATH mode: Jun 11 08:52:10
Entered DO_IMPORT mode: Jun 11 08:52:12
Entered DO_LABEL_ALLOC mode: Jun 11 08:52:16
Entered DO_RIBUPD mode: Jun 11 08:52:19
Entered Normal mode: Jun 11 08:52:23
Latest UPDATE sent: Jun 18 05:31:34

```

The following is sample output from show bgp process detail command with information on additional paths send and receive information:

```

BGP Process Information:
BGP is operating in STANDALONE mode
Autonomous System number format: ASDOT
Autonomous System: 100

```

show bgp process

```

Router ID: 22.22.22.22 (manually configured)
Default Cluster ID: 2.2.2.2 (manually configured)
Active Cluster IDs: 2.2.2.2
Fast external fallover enabled
Neighbor logging is enabled
Enforce first AS enabled
AS Path multipath-relax is enabled
Default local preference: 100
Default keepalive: 60
Graceful restart enabled
Restart time: 120
Stale path timeout time: 360
RIB purge timeout time: 600
Non-stop routing is enabled
Update delay: 120
Generic scan interval: 60

```

```
.....
```

```
.....
```

| | Allocated | Freed |
|-------------|-----------|-------|
| Prefixes: | 12 | 0 |
| Paths: | 60 | 0 |
| Path-elems: | 12 | 0 |

| | Number | Mem Used |
|-------------|--------|----------|
| Prefixes: | 12 | 1200 |
| Paths: | 60 | 3120 |
| Path-elems: | 12 | 624 |

Related Commands

| Command | Description |
|---|---|
| bgp bestpath as-path ignore, on page 56 | Sets the autonomous system path length to ignore when calculating preferred paths. |
| bgp bestpath compare-routerid, on page 58 | Compare identical routes received from external BGP (eBGP) peers during the best-path selection process and select the route with the lowest router ID. |
| bgp bestpath med always, on page 62 | Compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. |
| bgp bestpath med missing-as-worst, on page 66 | Assume paths with no MED attribute have the most undesirable MED value possible when performing path selection. |
| bgp cluster-id, on page 70 | Enables reflection of routes between route reflector clients using a BGP route reflector. |
| bgp cluster-id, on page 70 | Configure the cluster ID if the BGP cluster has more than one route reflector. |
| bgp default local-preference, on page 78 | Sets the default local preference value. |
| bgp redistribute-internal, on page 99 | Allows the redistribution of iBGP routes into an IGP such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF). |

| Command | Description |
|--|---|
| bgp router-id , on page 101 | Configures a fixed router ID for a BGP-speaking router. |
| default-metric (BGP) , on page 143 | Sets default metric values for the BGP. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| bgp scan-time , on page 103 | Configures scanning intervals. |
| timers bgp , on page 493 | Sets default BGP timers. |

show bgp regexp

To display routes matching the autonomous system path regular expression, use the **show bgp regexp** command in EXEC mode.

show bgp regexp *regular-expression*

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| <i>regular-expression</i> | Regular expression to match the BGP autonomous system paths. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each configured address family and subaddress family combination. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined in turn.

Use the **show bgp regexp** command to display all routes in the specified BGP table whose autonomous system path is matched by the specified regular expression.

**Note**

If the regular expression contains spaces and parentheses, it must be specified and surrounded by quotation marks.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp regexp** command:

```
RP/0/RSP0/CPU0:router# show bgp regexp "^3 "
BGP router identifier 10.0.0.5, local AS number 1
BGP main routing table version 64
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*>i172.20.17.121  10.0.101.2             100      0 3 2000 3000 i
*>i10.0.0.0        10.0.101.2             100      0 3 100 1000 i
*>i172.5.23.0/24  10.0.101.2             100      0 3 4 60 4378 i
```

This table describes the significant fields shown in the display.

Table 33: show bgp regexp Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|--------------|--|
| Origin codes | Origin of the path. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP address of a network entity. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp, on page 268 | Displays entries in the BGP routing table. |
| show bgp route-policy, on page 431 | Displays BGP information about networks that match an outbound route policy. |

```
show bgp regexp
```

show bgp route-policy

To display Border Gateway Protocol (BGP) information about networks that match an outbound route policy, use the **show bgp route-policy** command in EXEC mode.

```
show bgp route-policy route-policy-name []
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| <i>route-policy-name</i> | Name of a route policy. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

A route policy must be configured to use this command. When the **show bgp route-policy** command is entered, routes in the specified BGP table are compared with the specified route policy, and all routes passed by the route policy are displayed.

If a pass clause is encountered while the route policy is being applied to the route and the route policy processing completes without hitting a drop clause, the route is displayed. The route is not displayed if a drop clause is encountered, if the route policy processing completes without hitting a pass clause, or if the specified route policy does not exist.

The information displayed does not reflect modifications the policy might make to the route. To display such modifications, use the **show bgp policy** command.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp route-policy** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp route-policy p1

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 729
Dampening enabled
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, S stale
```

```

Origin codes: i - IGP, e - EGP, ? - incomplete
Network      Next Hop      Metric LocPrf Weight Path
* 10.13.0.0/16 192.168.40.24 0 1878 704 701 200 ?
* 10.16.0.0/16 192.168.40.24 0 1878 704 701 i

```

This table describes the significant fields shown in the display.

Table 34: show bgp route-policy Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|--------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 28 | Configures an aggregate entry in a BGP routing table. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor |
| route-policy | Configures a route policy. |

| Command | Description |
|--|---|
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp policy, on page 393 | Displays advertisements under a proposed policy. |

show bgp session-group

To display information about the Border Gateway Protocol (BGP) configuration for session groups, use the **show bgp session-group** command in EXEC mode.

```
show bgp session-group group-name {configuration [defaults] [nvgen] | inheritance | users}
```

Syntax Description

| | |
|----------------------|--|
| <i>group-name</i> | Name of the session family group to display. |
| configuration | (Optional) Displays the effective configuration for the session group, including any inherited configuration. |
| defaults | (Optional) Displays all configuration, including default configuration. |
| nvgen | (Optional) Displays output in the form of the show running-config command. If the defaults keyword also is specified, the output is not suitable for cutting and pasting into a configuration session. |
| inheritance | (Optional) Displays the session groups from which this session group inherits configuration. |
| users | (Optional) Display the session groups, neighbor groups, and neighbors that inherit configuration from this session group. |

Command Default

No default behavior or value

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp session-group** command with the *group-name* **configuration** argument and keyword to display the effective configuration of a session group, including any configuration inherited from other session groups through application of the **use** command. The source for each configured command is also displayed.

Use the **defaults** keyword to display the value of all configuration, including default configuration. Use the **nvgen** keyword to display configuration in the form of the **show running-config** command output. Output in this form is suitable for cutting and pasting into a configuration session.

Use the **show bgp session-group** command with the *group-name* **inheritance** argument and keyword to display the session groups from which the specified session group inherits configuration.

Use the **show bgp session-group** command with the *group-name* **users** argument and keyword to display the neighbors, neighbor groups, and session groups that inherit configuration from the specified session group.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

For the example shown here, the following configuration is used:

```
session-group group3
  advertisement-interval 5
  dmzlink-bw
  !
session-group group1
  use session-group group2
  update-source Loopback0
  !
session-group group2
  use session-group group3
  ebgp-multihop 2
```

The following example shows the **show bgp session-group** command with the **configuration** keyword:

```
RP/0/RSP0/CPU0:router# show bgp session-group group1 configuration

session-group group1
  advertisement-interval 5[s:group2 s:group3]
  ebgp-multihop 2 [s:group2]
  update-source Loopback0 []
  dmzlink-bandwidth [s:group2 s:group3]
```

The source of each command is shown to the right of the command. For example, **update-source** is configured directly on session group group1. The **dmzlink-bandwidth** command is inherited from session group group2, which in turn inherits it from session group group3.

The following example shows the **show bgp session-group** command with the **users** keyword:

```
RP/0/RSP0/CPU0:router# show bgp session-group group2 users

IPv4 Unicast:a:group1
```

The following example shows the **show bgp session-group** command with the **inheritance** keyword.

```
RP/0/RSP0/CPU0:router# show bgp session-group group1 inheritance

Session:s:group2 s:group3
```

The command output shows that the session group group1 directly uses the group2 session group. The group2 session group uses the group3 session group.

This table describes the significant fields shown in the display.

Table 35: show bgp session-group Field Descriptions

| Field | Description |
|-----------|--|
| [] | Configures the command directly on the specified session group. |
| s: | Indicates the name that follows is a session group. |
| a: | Indicates the name that follows is an address family group. |
| n: | Indicates the name that follows is a neighbor group. |
| [dflt] | Indicates the command is not explicitly configured or inherited, and the default value for the command is used. This field may be shown when the defaults keyword is specified. |
| <not set> | Indicates that the default is for the command to be disabled. This field may be shown when the defaults keyword is specified. |

Related Commands

| Command | Description |
|--|---|
| session-group, on page 264 | Configures a BGP session group. |
| show bgp neighbor-group, on page 349 | Displays information about the BGP configuration for neighbor groups. |
| show bgp neighbors, on page 353 | Displays information about BGP connections to neighbors. |

show bgp sessions

To display brief information about BGP neighbors, use the **show bgp sessions** command in EXEC mode.

show bgp sessions [not-established] [not-nsr-ready]

Syntax Description

| | |
|------------------------|---|
| not-established | (Optional) Displays all the neighbors that are not in established state |
| not-nsr-ready | (Optional) Displays all the neighbors that are not nonstop routing (NSR) ready. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show bgp sessions** command without a keyword provides brief information about all the BGP neighbors configured irrespective of the address family or VRF.

The **show bgp sessions** command with the **not-established** keyword shows BGP peers which are yet to establish their peering relationship.

The **show bgp sessions** command with the and **not-nsr-ready** keyword shows BGP peers which are yet to reach the nsr ready state.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp sessions** command in EXEC mode:

```
RP/0/RSP0/CPU0:router# show bgp sessions
Thu Jan 15 17:41:45.277 UTC
```

show bgp sessions

| Neighbor | VRF | Spk | AS | InQ | OutQ | NBRState | NSRState |
|--------------------------------|-------------------------|-----|-----|-----|------|-------------|-----------|
| 2.2.2.2 | default | 0 | 1 | 0 | 0 | Active | None |
| 10.0.101.1 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.2 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.3 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.4 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.5 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.6 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.7 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.8 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.0.101.9 | default | 0 | 1 | 0 | 0 | Established | NSR Ready |
| 10.11.12.2 | default | 0 | 100 | 0 | 0 | Established | NSR Ready |
| 90.0.0.2 | 900 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9000::1001 | 900 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 91.0.0.2 | 901 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9100::1001 | 901 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 92.0.0.2 | 902 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9200::1001 | 902 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 93.0.0.2 | 903 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9300::1001 | 903 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 94.0.0.2 | 904 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9400::1001 | 904 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 95.0.0.2 | 905 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9500::1001 | 905 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 96.0.0.2 | 906 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9600::1001 | 906 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 97.0.0.2 | 907 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9700::1001 | 907 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 98.0.0.2 | 908 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 9800::1001 | 908 | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 99.0.0.2 | 909 | 0 | 2 | 0 | 0 | Idle | None |
| 9900::1001 | 909 | 0 | 2 | 0 | 0 | Idle | None |
| 12.13.14.16 | red | 0 | 2 | 0 | 0 | Idle | None |
| 20.0.101.1 | red | 0 | 2 | 0 | 0 | Active | None |
| 1234:5678:9876::1111 | red | 0 | 3 | 0 | 0 | Idle | None |
| 2020::1002 | red | 0 | 2 | 0 | 0 | Established | NSR Ready |
| 1.2.3.4 | this-is-a-long-vrf-name | 0 | 5 | 0 | 0 | Idle | None |
| 1111:2222:3333:4444:5555::6789 | this-is-a-long-vrf-name | 0 | 7 | 0 | 0 | Idle | None |

The following is sample output from the **show bgp sessions** command with the **not-established** keyword:

```
RP/0/RSP0/CPU0:router# show bgp sessions not-established
Fri Jan 30 11:30:42.720 PST PDT
```

| Neighbor | VRF | Spk | AS | InQ | OutQ | NBRState | NSRState |
|---------------|---------|-----|-----|-----|------|----------|----------|
| 10.0.101.5 | default | 0 | 100 | 0 | 0 | Active | None |
| 2.2.2.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.101.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.102.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.103.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 4.4.4.2 | vrf1_1 | 0 | 304 | 0 | 0 | Idle | None |
| 2008:2:2:2::2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 11.16.1.2 | vrf2_1 | 0 | 302 | 0 | 0 | Idle | None |

The following is sample output from the **show bgp sessions** command with the **not-nsr-ready** keyword:

```
RP/0/RSP0/CPU0:router# show bgp sessions not-nsr-ready
Fri Jan 30 11:30:52.301 PST PDT
```

| Neighbor | VRF | Spk | AS | InQ | OutQ | NBRState | NSRState |
|------------|---------|-----|-----|-----|------|----------|----------|
| 10.0.101.5 | default | 0 | 100 | 0 | 0 | Active | None |
| 2.2.2.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.101.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.102.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 2.103.1.2 | vrf1_1 | 0 | 302 | 0 | 0 | Idle | None |
| 4.4.4.2 | vrf1_1 | 0 | 304 | 0 | 0 | Idle | None |

```

2008:2:2:2::2   vrf1_1           0  302   0   0  Idle   None
11.16.1.2      vrf2_1           0  302   0   0  Idle   None

```

This table describes the significant fields shown in the display.

Table 36: show bgp sessions Field Descriptions

| Field | Description |
|----------|---|
| Neighbor | Displays neighbor IP address. |
| VRF | Displays information about the VRF. |
| Spk | Speaker process that is responsible for the neighbor. Always 0. |
| AS | Autonomous system. |
| InQ | Number of messages from a neighbor waiting to be processed. |
| OutQ | Number of messages waiting to be sent to a neighbor. |
| NBRState | State of the Border Gateway Protocol (BGP) neighbor sessions. |
| NSRState | State of the Border Gateway Protocol (BGP) nonstop routing (NSR). |

Related Commands

| Command | Description |
|---|--|
| show bgp neighbors, on page 353 | Displays information about Border Gateway Protocol (BGP) connections to neighbors. |

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6 {unicast| multicast| all|
labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vrf {vrf-name|
all}| [ipv4 {unicast| labeled-unicast}| ipv6 unicast]| vpnv6 unicast] summary
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Use the **show bgp summary** command to display a summary of the neighbors for which the specified address family and subaddress family are enabled. If the neighbor does not have the specified address family and subaddress family enabled, it is not included in the output of the **show** command. If the **all** keyword is specified for the address family or subaddress family, a summary for each combination of address family and subaddress family is displayed in turn.

The table versions shown in the output (RcvTblVer, bRIB/RIB, SendTblVer, and TblVer) are specific to the specified address family and subaddress family. All other information is global.

The table versions provide an indication of whether BGP is up to date with all work for the specified address family and subaddress family.

- bRIB/RIB < RcvTblVer—Some received routes have not yet been considered for installation in the global routing table.
- TblVer < SendTblVer—Some received routes have been installed in the global routing table but have not yet been considered for advertisement to this neighbor.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp summary** command:

```
RP/0/RSP0/CPU0:router#show bgp summary
```

show bgp summary

```

BGP router identifier 10.0.0.0, local AS number 2
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

```

BGP is operating in STANDALONE mode.

```

Process          RecvTblVer    bRIB/RIB    LabelVer    ImportVer    SendTblVer
Speaker          1             0           1           1           0

Neighbor        Spk   AS  MsgRcvd  MsgSent    TblVer    InQ  OutQ  Up/Down  St/PfxRcd
10.0.101.0      0     2     0        0         0         0    0    00:00:00 Idle
10.0.101.1      0     2     0        0         0         0    0    00:00:00 Idle

```

This table describes the significant fields shown in the display.

Table 37: show bgp summary Field Descriptions

| Field | Description |
|--------------------------------|--|
| BGP router identifier | IP address of the router. |
| local AS number | Autonomous system number set by the router bgp, on page 253 command. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | Specifies BGP is operating in standalone mode. |
| Process | BGP process. |

| Field | Description |
|---|--|
| RecvTblVer | Last version used in the BGP database for received routes. |
| bRIB/RIB | Last version of the local BGP database that was injected into the main routing table. |
| LabelVer | Label version used in the BGP database for label allocation. |
| ImportVer | Last version of the local BGP database for importing routes. |
| SendTblVer | Latest version of the local BGP database that is ready to be advertised to neighbors. |
| Some configured eBGP neighbors do not have any policy | Some external neighbors exist that do not have both an inbound and outbound policy configured for every address family, using the route-policy (BGP) command. In this case, no prefixes are accepted and advertised to those neighbors. |
| Neighbor | IP address of a neighbor. |
| Spr | Speaker process that is responsible for the neighbor. Always 0. |
| AS | Autonomous system. |
| MsgRcvd | Number of BGP messages received from a neighbor. |
| MsgSent | Number of BGP messages sent to a neighbor. |
| TblVer | Last version of the BGP database that was sent to a neighbor. |
| InQ | Number of messages from a neighbor waiting to be processed. |
| OutQ | Number of messages waiting to be sent to a neighbor. |
| Up/Down | Length of time in (hh:mm:ss) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established. |

| Field | Description |
|-----------|--|
| St/PfxRcd | <p>If the BGP session is not established, the current state of the session. If the session is established, the number of prefixes the router has received from the neighbor.</p> <p>If the number of prefixes received exceeds the maximum allowed (as set by the maximum-prefix command), "(PfxRcd)" appears.</p> <p>If the connection has been shut down using the shutdown command, "(Admin)" appears.</p> <p>If the neighbor is external and it does not have an inbound and outbound policy configured for every address family, an exclamation mark (!) is inserted at the end of the state when using the route-policy (BGP) command.</p> <p>If the connection has been shut down due to out of memory (OOM), "(OOM)" appears.</p> |

Related Commands

| Command | Description |
|---|--|
| route-policy (BGP), on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |

show bgp summary nsr

To display the summary of Border Gateway Protocol (BGP) neighbor state and nonstop routing (NSR) state information, use the **show bgp summary nsr** command in EXEC mode.

```
show bgp summary [ipv4 {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6 {unicast| multicast|
all| labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vrf {vrf-name|
all} [ipv4 {unicast| labeled-unicast}| ipv6 unicast]| vpnv6 unicast] nsr [standby]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| mdt | (Optional) Specifies multicast distribution tree (MDT) address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| vpnv6 unicast | (Optional) Specifies VPNv6 unicast address families. |
| standby | Displays information about the standby card. |

Command Default If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples The following is sample output from the **show bgp summary nsr** command:

```
RP/0/RSP0/CPU0:router# show bgp summary nsr

BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 11034 (Reached)
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

node0_1_CPU0          Speaker

Entered mode Standby Ready           : Feb  3 14:22:00
Entered mode TCP NSR Setup           : Feb  3 14:22:00
Entered mode TCP NSR Setup Done      : Feb  3 14:22:01
Entered mode TCP Initial Sync        : Feb  3 14:22:01
Entered mode TCP Initial Sync Done   : Feb  3 14:22:44
Entered mode FPBSN processing done   : Feb  3 14:22:44
Entered mode Update processing done   : Feb  3 14:22:44
Entered mode BGP Initial Sync        : Feb  3 14:22:44
Entered mode BGP Initial Sync done   : Feb  3 14:22:49
Entered mode NSR Ready                : Feb  3 14:22:49

Current BGP NSR state - NSR Ready achieved at: Feb  3 14:22:49
NSR State READY notified to Redcon at: Feb  4 07:44:43

Process      RcvTblVer  bRIB/RIB  LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker      13037      13037     13037     13037     13037       13037
```

```
Neighbor      Spk   AS   TblVer  SyncVer  AckVer  NBRState   NSRState
2.2.2.2       0    302  13037   13037    13037  Established NSR Ready
10.0.101.5    0    100  13037   13037    13037  Established NSR Ready
```

The following example shows sample output from the **show bgp summary nsr** command with the **standby** keyword:

```
RP/0/RSP0/CPU0:router# show bgp summary nsr standby
```

```
BGP router identifier 10.1.0.1, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 13037
BGP NSR Initial initsync version 0 (Not Reached)
BGP scan interval 60 secs
```

BGP is operating in STANDALONE mode.

```
node0_0_CPU0      Speaker
```

```
Entered mode Standby Ready           : Feb  3 14:22:03
Entered mode TCP Replication         : Feb  3 14:22:03
Entered mode TCP Init Sync Done      : Feb  3 14:22:47
Entered mode NSR Ready                : Feb  3 14:22:52
```

```
Process      RcvTblVer  bRIB/RIB  LabelVer  ImportVer  SendTblVer  StandbyVer
Speaker      13037      0          0          13037      0            0
```

```
Neighbor      Spk   AS   TblVer  SyncVer  AckVer  NBRState   NSRState
2.2.2.2       0    302  13037   0        1       Established NSR Ready
10.0.101.5    0    100  13037   0        1       Established NSR Ready
```

This table describes the significant fields shown in the display.

Table 38: show bgp summary nsr Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | IP address of the router. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| Non-stop routing | State of the Nonstop routing. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | Specifies BGP is operating in standalone mode. |

| Field | Description |
|--------------|--|
| Entered mode | The successive transition of various states of TCP and BGP, leading to the NSR ready state. Note This is used for monitoring and debugging purposes. |
| SyncVer | The version which has synced to standby for this neighbor. |
| AckVer | The version which the neighbor has acknowledge. |
| NBRState | State of the BGP neighbor. |
| NSRState | Neighbor NSR state. |

Related Commands

| Command | Description |
|---|---|
| nsr (BGP), on page 210 | Activates Border Gateway Protocol (BGP) nonstop routing (NSR) |
| show bgp nsr, on page 386 | Displays Border Gateway Protocol (BGP) nonstop routing (NSR) information. |

show bgp truncated-communities

To display routes in the Border Gateway Protocol (BGP) routing table for which inbound policy or aggregation has exceeded the maximum number of communities that may be attached, use the **show bgp truncated-communities** command in EXEC mode.

show bgptruncated-communities

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) For subaddress families, specifies prefixes for all subaddress families. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| all | (Optional) For address family, specifies prefixes for all address families. |
| vpn4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| vrf-name | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

BGP contains a separate routing table for each address family and subaddress family combination that has been configured. The address family and subaddress family options specify the routing table to be examined. If the **all** keyword is specified for the address family or subaddress family, each matching routing table is examined.

Use the **show bgp truncated-communities** command to display those routes in the specified BGP routing table in which the buffers used to store communities or extended communities have overflowed. An overflow occurs if an attempt is made to associate more communities or extended communities with the route than fits in a BGP update message. This can happen due to modification of communities or extended communities during aggregation or when inbound policy is applied.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |

Examples The following is sample output from the **show bgp truncated-communities** command:

```
RP/0/RSP0/CPU0:router# show bgp truncated-communities

BGP router identifier 172.20.1.1, local AS number 1820
BGP main routing table version 3042
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
*  10.13.0.0/16   192.168.40.24           0 1878 704 701 200 ?
```

```
*> 10.16.0.0/16      192.168.40.24      0 1878 704 701 i
```

This table describes the significant fields shown in the display.

Table 39: show bgp truncated-communities Field Descriptions

| Field | Description |
|--------------------------------|---|
| BGP router identifier | BGP Identifier for the local system. |
| local AS number | Autonomous system number for the local system. |
| BGP main routing table version | Last version of the BGP database that was installed into the main routing table. |
| Dampening enabled | Displayed if dampening is enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| Status codes | <p>Status of the table entry. The status is displayed as a three-character field at the beginning of each line in the table. The first character may be (in order of precedence):</p> <ul style="list-style-type: none"> S—Path is stale, indicating that a graceful restart is in progress with the peer from which the route was learned. s—Path is more specific than a locally sourced aggregate route and has been suppressed. *—Path is valid. <p>The second character may be (in order of precedence):</p> <ul style="list-style-type: none"> >—Path is the best path to use for that network. d—Path is dampened. h—Path is a history entry, representing a route that is currently withdrawn, but that is being maintained to preserve dampening information. Such routes should never be marked as valid. <p>The third character may be:</p> <ul style="list-style-type: none"> i—Path was learned by an internal BGP (iBGP) session. |

| Field | Description |
|--------------|---|
| Origin codes | Origin of the path. The origin code is displayed at the end of each line in the table. It can be one of the following values: i—Path originated from an Interior Gateway Protocol (IGP) and was advertised with a network or aggregate-address command. e—Path originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP. |
| Network | IP prefix and prefix length for a network. |
| Next Hop | IP address of the next system that is used when a packet is forwarded to the destination network. An entry of 0.0.0.0 indicates that the router has a non-BGP route to this network. |
| Metric | Value of the interautonomous system metric, otherwise known as the Multi Exit Discriminator (MED) metric. |
| LocPrf | Local preference value. This is used to determine the preferred exit point from the local autonomous system. It is propagated throughout the local autonomous system. |
| Weight | Path weight. Weight is used in choosing the preferred path to a route. It is not advertised to any neighbor. |
| Path | Autonomous system path to the destination network. At the end of the path is the origin code for the path. |

Related Commands

| Command | Description |
|--|---|
| aggregate-address , on page 28 | Creates an aggregate entry in a BGP routing table. |
| network (BGP) , on page 196 | Specifies a local network that the BGP routing process should originate and advertise to its neighbors. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |

| Command | Description |
|--|---|
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp , on page 268 | Displays entries in the BGP routing table. |

show bgp update-group

To display Border Gateway Protocol (BGP) information for update groups, use the **show bgp update-group** command in EXEC mode.

```
show bgp [ipv4 {unicast| multicast| labeled-unicast| all| tunnel| mdt}| ipv6 {unicast| multicast| all|
labeled-unicast}| all {unicast| multicast| all| labeled-unicast| mdt| tunnel}| vpnv4 unicast| vrf {vrf-name|
all}| [ipv4 {unicast| labeled-unicast}| ipv6 unicast]| vpnv6 unicast] update-group [neighbor ip-address|
process-id.index [summary| performance-statistics]]
```

Syntax Description

| | |
|---|---|
| ipv4 | (Optional) Specifies IP Version 4 update groups. |
| unicast | (Optional) Specifies unicast update groups. |
| multicast | (Optional) Specifies multicast update groups. |
| labeled-unicast | (Optional) Specifies labeled unicast address prefixes. |
| all | (Optional) Displays both unicast and multicast update groups. |
| tunnel | (Optional) Specifies tunnel address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 update groups. |
| all | (Optional) Displays both IP Version 4 and IP Version 6 update groups. |
| vpnv4 unicast | (Optional) Specifies VPNv4 unicast address families. |
| rd rd-address | (Optional) Displays routes with a specific route distinguisher. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of a VRF. |
| all | (Optional) For VRF, specifies all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) For VRF, specifies IPv4 unicast or labeled-unicast address families. |
| ipv6 unicast | (Optional) For VRF, specifies IPv6 unicast address families. |
| neighbor ip-address | (Optional) Specifies information on an update group for a specific neighbor. |
| <i>process-id.index</i> | (Optional) Update group index. Process ID range is 0 to 254. Index range is 0 to 4294967295. The <i>process id.index</i> argument is specified as follows: process ID (dot) index. In standalone mode, the process ID is always 0. |

| | |
|-------------------------------|--|
| summary | (Optional) Specifies summary of update group members. |
| performance-statistics | (Optional) Specifies performance information about the updates generated for the update group. |

Command Default

If no address family or subaddress family is specified, the default address family and subaddress family specified using the **set default-afi** and **set default-safi** commands are used.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 5.1.1 | The command output was modified to include the status of advertised permanent paths. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

The **set default-afi** command is used to specify the default address family for the session, and the **set default-safi** command is used to specify the default subaddress family for the session. See the *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for detailed information and syntax for the **set default-afi** and **set default-safi** commands. If you do not specify a default address family, the default address family is IPv4. If you do not specify a default subaddress family, the default subaddress family is unicast.

Every BGP neighbor is automatically assigned to an update group for each address family that is enabled on the neighbor. Neighbors that have similar outbound policy, such that they are sent the same updates, are placed in the same update group.

Use the **show bgp update-group** command to display the update groups and a list of the neighbors that belong to the update group.

Use the **show bgp update-group neighbor** command to display details about the update group to which a neighbor belongs for the specified address family.

Use the **summary** keyword to display a summary of the neighbors belonging to the specified update group. The display format is the same as for the [show bgp summary](#), on page 442 command.

Use the **performance-statistics** keyword to display information about the number of prefixes processed and the time taken to generate updates for the specified update group.

**Note**

Update group indexes are not necessarily persistent over a process restart. If a BGP process restarts, the index of the update group to which a particular neighbor is assigned may be different, though the set of neighbors belonging to the update group is the same.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the **show bgp update-group** command:

```
RP/0/RSP0/CPU0:router# show bgp update-group

Update group for IPv4 Unicast, index 0.1:
  Attributes:
    Internal
    Common admin
    Send communities
    Send extended communities
    Minimum advertisement interval: 300
    Update group desynchronized: 0
    Sub-groups merged: 0
    Messages formatted: 0, replicated: 0
    Neighbors not in any sub-group:
      10.0.101.1
```

This table describes the significant fields shown in the display.

Table 40: show bgp update-group Field Descriptions

| Field | Description |
|------------------|--|
| Update group for | Address family to which updates in this update group apply. |
| index | Update group index. |
| Attributes | Attributes common to all members of the update group. |
| Unsuppress map | Unsuppress route map used to selectively unsuppress more specific routes of locally generated aggregates for members of this update group. |

| Field | Description |
|--------------------------------|---|
| Outbound policy | Route policy applied to outbound updates generated for members of this update group. |
| Internal | Members of the update group are internal peers. |
| ORF Receive enabled | Members of this update group are capable of receiving an outbound route filter. |
| Route Reflector Client | Local system is acting as a route reflector for members of this update group. |
| Remove private AS numbers | Members of this update group have private AS numbers stripped from outbound updates. |
| Next-hop-self enabled | Next- Next hop for members of the update group is set to the local router. |
| Directly connected IPv6 EBGp | Members of this update group are directly connected external BGP IPv6-based peers. |
| Configured Local AS | Local autonomous system (AS) used for members of this update group. |
| Common admin | Peers in this update group are under common administration (internal or confederation peers). |
| Send communities | Communities are sent to neighbors in this update group. |
| Send extended communities | Extended communities is sent to neighbors in this update group. |
| Minimum advertisement interval | Minimum advertisement interval for members of this update group. |
| replicated | Number of update messages replicated for this update group. |
| Messages formatted | Number of update messages generated for this update group. |
| Neighbors in this update group | List of neighbors that use this update group for the given address family. |
| Update group desynchronized | Number of times an update group has been split to accommodate the slower peer. This option is disabled. |
| Sub-groups merged | Number of times an update group has been split and merged. |

| Field | Description |
|--------------------------------|--|
| Neighbors not in any sub-group | BGP neighbor that does not belong to any subgroup. |

The following is sample output from the **show bgp update-group** command with the **ipv4**, **unicast**, and **summary** keywords and the *process id.index* argument:

```
RP/0/RSP0/CPU0:router# show bgp ipv4 unicast update-group 0.1 summary

BGP router identifier 10.140.140.1, local AS number 1.1
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000
BGP main routing table version 1
BGP scan interval 60 secs

BGP is operating in STANDALONE mode.

Process          RecvTblVer    bRIB/RIB  LabelVer  ImportVer  SendTblVer
Speaker          1             0          1          1           0

Neighbor         Spr    AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down  St/PfxRcd
172.25.11.8      0     1     0        0         0     0    0 00:00:00 Idle
```

This is sample output from the **show bgp ipv4 unicast update-group** command showing the status of advertised permanent paths:

```
RP/0/RSP0/CPU0:router# show bgp ipv4 unicast update-group
Update group for IPv4 Unicast, index 0.2:
  Attributes:
    Neighbor sessions are IPv4
    Outbound policy: PASS
    Internal
    Common admin
    First neighbor AS: 30813
    Send communities
    Send extended communities
    Next-hop-self enabled
    4-byte AS capable
    Non-labeled address-family capable
    Advertise Permanent-Network capable
    Send AIGP
    Minimum advertisement interval: 0 secs
    Update group desynchronized: 0
    Sub-groups merged: 4
    Number of refresh subgroups: 0
    Messages formatted: 42, replicated: 68
    Neighbors not in any sub-group:
      100.12.13.3    100.13.13.3
```

This table describes the significant fields shown in the display.

Table 41: show bgp ipv4 unicast update-group Field Descriptions

| Field | Description |
|-----------------------|---------------------------|
| BGP router identifier | IP address of the router. |

| Field | Description |
|--------------------------------|---|
| local AS number | Autonomous system number set by the router bgp , on page 253 command. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| BGP generic scan interval | Interval (in seconds) between scans of the BGP table by a generic scanner. |
| BGP table state | State of the BGP database. |
| Table ID | BGP database identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| Dampening enabled | Displayed if dampening has been enabled for the routes in this BGP routing table. |
| BGP scan interval | Interval (in seconds) between scans of the BGP table specified by the address family and subaddress family. |
| BGP is operating in | BGP is operating in standalone mode. |
| Process | BGP process. |
| RecvTblVer | Last version used in the BGP database for received routes. |
| bRIB/RIB | Last version of the local BGP database that was injected into the main routing table. |
| LabelVer | Label version used in the BGP database for label allocation. |
| ImportVer | Last version of the local BGP database for importing routes. |
| SendTblVer | Latest version of the local BGP database that is ready to be advertised to neighbors. |

| Field | Description |
|---|---|
| Some configured eBGP neighbors do not have any policy | Some external neighbors that exist do not have both an inbound and outbound policy configured for every address family, using the route-policy (BGP) command. In this case, no prefixes are accepted or advertised to those neighbors. |
| Neighbor | IP address of a neighbor. |
| Spr | Speaker process that is responsible for the neighbor. Always 0. |
| AS | Autonomous system. |
| MsgRcvd | Number of BGP messages received from a neighbor. |
| MsgSent | Number of BGP messages sent to a neighbor. |
| TblVer | Last version of the BGP database that was sent to a neighbor. |
| InQ | Number of messages from a neighbor waiting to be processed. |
| OutQ | Number of messages waiting to be sent to a neighbor. |
| Up/Down | Length of time (in hh:mm:s) that the BGP session has been in Established state, or the time since the session left Established state, if it is not established. |
| St/PfxRcd | <p>If the BGP session is not established, the current state of the session. If the session is established, the number of prefixes the router has received from the neighbor.</p> <p>If the number of prefixes received exceeds the maximum allowed (as set by the maximum-prefix command), "(PfxRcd)" appears.</p> <p>If the connection has been shut down using the shutdown command, "(Admin)" appears.</p> <p>If the neighbor is external and it does not have an inbound and outbound policy configured for every address family, an exclamation mark (!) is inserted at the end of the state when using the route-policy (BGP) command.</p> |

Related Commands

| Command | Description |
|--|--|
| maximum-prefix (BGP) , on page 185 | Limits the number of prefixes that can be received from a neighbor. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |
| set default-afi | Sets the default Address Family Identifier (AFI) for the current session. |
| set default-safi | Sets the default Subaddress Family Identifier (SAFI) for the current session. |
| show bgp summary , on page 442 | Displays the status of all BGP connections. |
| shutdown (BGP) , on page 473 | Disables a neighbor without removing its configuration. |

show bgp vrf imported-routes

To display Border Gateway Protocol (BGP) information for routes imported into specified VPN routing and forwarding (VRF) instances, use the **show bgp vrf imported-routes** command in EXEC mode.

show bgp vrf{vrf| *vrf-name*} **imported-routes**

Syntax Description

| | |
|---|---|
| <i>vrf-name</i> | Displays imported routes for a specific VRF. |
| all | Displays imported routes for all VRFs. |
| ipv4 { unicast labeled-unicast } | (Optional) Specifies IP Version 4 unicast or labeled-unicast imported routes. |
| ipv6 unicast | (Optional) Specifies IP Version 6 unicast imported routes. |
| vrf <i>source-vrf-name</i> | (Optional) Displays routes imported from the specified source VRF. |
| neighbor <i>neighbor-address</i> | (Optional) Displays preview advertisements for a specified neighbor. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show bgp vrf imported-routes** command to display all paths imported into a specified VRF from the default VRF. Use the **neighbor** *neighbor-address* keyword and argument to display all imported paths and which paths were learned from the specified neighbor. Use the **vrf** *source-vrf-name* keyword and argument to display all imported routes that belong to the specified source VRF. The **neighbor** *neighbor-address* and **vrf** *source-vrf-name* cannot coexist.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |

Examples

The following is sample output from the `show bgp vrf imported-routes` command:

```
RP/0/RSP0/CPU0:router# show bgp vrf vrf-1 ipv6 unicast imported-routes

BGP VRF one, state: Active BGP
BGP Route Distinguisher: 100:222
VRF ID: 0x60000001
BGP router identifier 10.2.0.1, local AS number 100
BGP table state: Active
Table ID: 0xe0800001
BGP main routing table version 41534

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Neighbor      Route Distinguisher   Source VRF
*>i1234:1052::/32    10.1.0.1      100:111                default
*>i2008:1:1:1::/112  10.1.0.1      100:111                default
*>i2008:111:1:1::1/128
                    10.1.0.1      100:111                default

Processed 3 prefixes, 3 paths
```

This table describes the significant fields shown in the display.

Table 42: show bgp vrf imported-routes Field Descriptions

| Field | Description |
|--------------------------|--|
| BGP VRF | VRF name. |
| state | State of the VRF. |
| BGP Route Distinguisher: | Unique identifier for the BGP routing instance. |
| VRF Id | VRF identifier. |
| BGP router identifier | IP address of the router. |
| local AS number | Autonomous system number set by the router bgp, on page 253 command. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |

| Field | Description |
|--------------------------------|---|
| BGP table state | State of the BGP database. |
| Table ID | Table identifier. |
| BGP main routing table version | Last version of the BGP database that was injected into the main routing table. |
| Network | Network address. |
| Neighbor | IP address of a neighbor. |
| Route Distinguisher | Unique identifier for the routing instance. |
| Source VRF | Source VRF for the imported route. |

show protocols (BGP)

To display information about the Border Gateway Protocol (BGP) instances running on the router, use the **show protocols** command in EXEC mode and specify either the **bgp** or **all** keyword.

```
show protocols [ipv4|ipv6|afi-all] [all|protocol]
```

Syntax Description

| | |
|-----------------|---|
| ipv4 | (Optional) Specifies the IP Version 4 address family. |
| ipv6 | (Optional) Specifies the IP Version 6 address family. |
| afi-all | (Optional) Specifies all address families. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are bgp , isis , rip , eigrp , and ospf . For the IPv6 address family, the options are bgp , eigrp , isis , and ospfv3 . |

Command Default

Default is IPv4.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected. For BGP, the command output lists the protocol ID, peers with elapsed time since last reset, and miscellaneous information, such as external and internal local distances and sourced routes.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| rib | read |

Examples

The following example shows the display for the **show protocols** command using the **bgp** keyword:

```
RP/0/RSP0/CPU0:router# show protocols bgp
Routing Protocol "BGP 40"
Address Family IPv4 Unicast:
  Distance: external 20 internal 200 local 200
  Sourced Networks:
    10.100.0.0/16 backdoor
    10.100.1.0/24
    10.100.2.0/24
  Routing Information Sources:
    Neighbor      State/Last update received
    10.5.0.2      Idle
    10.9.0.3      Idle
```

This table describes the significant fields shown in the display.

Table 43: show protocols (BGP) Field Descriptions

| Field | Description |
|--------------------|--|
| Routing Protocol: | Identifies BGP as the running protocol and displays the BGP AS number. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Address Family | Specifies the address family. This can be IPv4 Unicast, IPv4 Multicast, or IPv6 Unicast. |
| Distance: external | Specifies the distance BGP sets when installing eBGP routes into the RIB. eBGP routes are routes received from eBGP peers. The RIB uses the distance as a tiebreaker when several protocols install a route for the same prefix. |
| Distance: internal | Specifies the distance BGP sets for routes received from iBGP peers. |

| Field | Description |
|-----------------------------|--|
| Distance: local | Specifies the distance BGP sets for locally generated aggregates and backdoor routes. |
| Sourced Networks | List of locally sourced networks. These are networks sourced using the network command. |
| Routing information Sources | List of configured BGP neighbors. |
| Neighbor | Address of a BGP neighbor. |
| State/Last update received | State of each neighbor and the time since the last update was received from the neighbor if it is established. |

show svd role

To display selective VRF download (SVD) role information, use the **show svd role** command in EXEC mode.

show svd role

Syntax Description This command has no keywords or arguments.

Command Default None.

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show svd role** command output displays name of the line card and role for each address-family in a table.

| Task ID | Task ID | Operation |
|---------|-------------|-----------|
| | ip-services | read |

Examples This example displays the different nodes in a line card and corresponding IPv4, and IPv6 SVD role information:

```
RP/0/RSP0/CPU0:router#show svd role
Thu Mar 10 10:45:17.886 PST
Node Name      IPv4 Role      IPv6 Role
-----
0/1/CPU0      Core Facing    Not Interested
0/2/CPU0      Core Facing    Core Facing
0/4/CPU0      Standard       Standard
0/5/CPU0      Standard       Standard
```

| Related Commands | Command | Description |
|------------------|--|--|
| | selective-vrf-download disable , on page 256 | Disables selective VRF Download (SVD). |

| Command | Description |
|---|--|
| show svd state, on page 472 | Displays Selective VRF Download (SVD) state information. |

show svd state

To display selective VRF download (SVD) state information, use the **show svd state** command in EXEC mode.

show svd state

Syntax Description This command has no keywords or arguments.

Command Default None.

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|-------------|-----------|
| | ip-services | read |

Examples This example shows the SVD configuration state and the SVD operational state in a line card:

```
RP/0/RSP0/CPU0:router#show svd state
Thu Mar 10 10:45:32.184 PST
Selective VRF Download (SVD) Feature State:
SVD Configuration State      Enabled
SVD Operational State        Enabled
```

| Related Commands | Command | Description |
|------------------|---|---|
| | selective-vrf-download disable, on page 256 | Disables selective VRF Download (SVD). |
| | show svd role, on page 470 | Displays Selective VRF Download (SVD) role information. |

shutdown (BGP)

To disable a neighbor without removing its configuration, use the **shutdown** command in an appropriate configuration mode. To re-enable the neighbor and reestablish a Border Gateway Protocol (BGP) session, use the **no** form of this command.

shutdown [inheritance-disable]

no shutdown [inheritance-disable]

Syntax Description

| | |
|----------------------------|---|
| inheritance-disable | (Optional) Overrides the value of a shutdown command inherited from a neighbor group or session group. |
|----------------------------|---|

Command Default

Neighbors are not shutdown.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **shutdown** command to terminate any active session for the specified neighbor and remove all associated routing information. Use of the **shutdown** command with a neighbor group or session group may suddenly terminate a large number of BGP neighbor sessions because all neighbors using the neighbor group or session group may be affected.

Use the **show bgp summary** command to display a summary of BGP neighbors. Neighbors that are idle due to the **shutdown** command are displayed with the "Idle (Admin)" state.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows that any active session for neighbor 192.168.40.24 is disabled:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# shutdown
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

In the following example, the session remains active for neighbor 192.168.40.24 because the inherited **shutdown** command has been overridden:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group group1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# shutdown
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# shutdown inheritance-disable
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| show bgp summary, on page 442 | Displays the status of all BGP connections. |

shutdown (rpki-server)

To shutdown RPKI cache-server, use the **shutdown** command in rpki-server configuration mode. To set that the RPKI cache be active, use the **no** form of this command.

shutdown

no shutdown

This command has no keywords or arguments.

Command Default RPKI cache is active.

Command Modes RPKI server configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 4.2.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|----------------|----------------|------------------|
| | bgp | read, write |

Examples This command shows how to configure no shutdown of the RPKi cache configuration after other RPKI cache parameters are configured:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#password rpki-ssh-pass
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#preference 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#purge-time 30
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#refresh-time 30
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#response-time 30
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#no shutdown
```

signalling disable

To disable BGP or LDP signaling protocol to neighbors, use the **signalling disable** command in neighbor address family (l2vpn vpls-vpws) configuration mode. To restore the system to its default condition, use the **no** form of this command.

signalling {bgp|ldp} disable

| Syntax Description | | |
|--------------------|--|--|
| bgp | Selects BGP signaling protocol to disable. | |
| ldp | Selects LDP signaling protocol to disable. | |

Command Default Both BGP and LDP signaling protocols are enabled.

Command Modes Neighbor address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.9.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | bgp | read, write |

Examples The following example shows how to disable BGP signaling protocol for neighbor 10.2.3.4:

```
RP/0/RSP0/CPU0:router (config) #router bgp 100
RP/0/RSP0/CPU0:router (config-bgp) #neighbor 10.2.3.4
RP/0/RSP0/CPU0:router (config-bgp-nbr) #address-family l2vpn vpls-vpws
RP/0/RSP0/CPU0:router (config-bgp-nbr-af) #signalling bgp disable
```

site-of-origin (BGP)

To attach a site-of-origin extended community attribute to each route received from the specified peer, use the **site-of-origin** command in VRF neighbor address family configuration mode. To restore the system to its default condition, use the **no** form of this command.

site-of-origin [*as-number:nn* | *ip-address:nn*]

Syntax Description

as-number:nn

- *as-number*— Autonomous system (AS) number.
 - Range for 2-byte Autonomous system number is 1 to 65535.
 - Range for 4-byte Autonomous system number in asplain format is 1 to 4294967295.
 - Range for 4-byte Autonomous system number in asdot format is 1.0 to 65535.6553.
- *nn* —32-bit number

ip-address:nn

IP address.

- *ip-address* —32-bit IP address
- *nn* —16-bit number

Command Default

No default behavior or values

Command Modes

VRF neighbor address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When routes are advertised to the peer, routes whose extended communities list contain the site of origin (SoO) are filtered out and not advertised to the peer. Site-of-origin uniquely identifies the site from which the provide edge (PE) router learned routes, thus filtering based on the extended community helps prevent transient routing loops from occurring in complex and mixed network topologies.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure SoO filtering:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf_A
RP/0/RSP0/CPU0:router(config-bgp-vrf)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# remote-as 10
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-nbr-af)# site-of-origin 10.0.01:20
```

socket receive-buffer-size

To set the size of the receive buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket receive-buffer-size** command in an appropriate configuration mode. To set the size of the receive buffers to the default size, use the **no** form of this command.

socket receive-buffer-size *socket-size* [*bgp-size*]

no socket receive-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

| | |
|--------------------|---|
| <i>socket-size</i> | Size (in bytes) of the receive-side socket buffers. Range is 512 to 131072. |
| <i>bgp-size</i> | (Optional) Size (in bytes) of the receive buffers in BGP. Range is 512 to 131072. |

Command Default

socket-size : 32,768 bytes

bgp-size : 4,032 bytes

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **socket receive-buffer-size** command to increase the buffer size when receiving updates from a neighbor. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.



Note

Increasing the socket buffer size uses more memory only when more messages are waiting to be processed by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **receive-buffer-size** command on individual neighbors to change the values set by the **socket receive-buffer-size** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the receive buffer sizes for all neighbors to 65,536 bytes for the socket buffer and 8192 bytes for the BGP buffer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# socket receive-buffer-size 65536 8192
```

Related Commands

| Command | Description |
|--|--|
| receive-buffer-size, on page 229 | Sets the size of the receive buffers for a BGP neighbor. |
| socket send-buffer-size, on page 481 | Sets the size of the send buffers for all BGP neighbors. |

socket send-buffer-size

To set the size of the send buffers for all Border Gateway Protocol (BGP) neighbors, use the **socket send-buffer-size** command in an appropriate configuration mode. To set the size of the send buffers to the default size, use the **no** form of this command.

socket send-buffer-size *socket-size* [*bgp-size*]

no socket send-buffer-size [*socket-size*] [*bgp-size*]

Syntax Description

| | |
|--------------------|---|
| <i>socket-size</i> | Size (in bytes) of the send-side socket buffers. Range is 4096 to 131072. |
| <i>bgp-size</i> | (Optional) Size (in bytes) of the send buffers in BGP. Range is 4096 to 131072. |

Command Default

socket-size : 10240 bytes

bgp-size : 4096 bytes

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **socket send-buffer-size** command to increase the buffer size when sending updates to neighbors. Using larger buffers can improve convergence time because the software can process more packets simultaneously. However, allocating larger buffers uses more memory on your router.



Note

Increasing the socket buffer size uses more memory only when more messages are waiting to be sent by the software. In contrast, increasing the BGP buffer size uses extra memory indefinitely.

Use the **send-buffer-size** command on individual neighbors to change the values set by the **socket send-buffer-size** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to set the send buffer sizes for all neighbors to 8192 bytes for the socket buffer and the BGP buffer:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# socket send-buffer-size 8192 8192
```

Related Commands

| Command | Description |
|---|---|
| send-buffer-size, on page 258 | Sets the size of the send buffers for a BGP neighbor. |
| socket receive-buffer-size, on page 479 | Sets the size of the receive buffers for all BGP neighbors. |

soft-reconfiguration inbound

To configure the software to store updates received from a neighbor, use the **soft-reconfiguration inbound** command in an appropriate configuration mode. To disable storing received updates, use the **no** form of this command.

soft-reconfiguration inbound [**always**| **inheritance-disable**]

no soft-reconfiguration inbound [**always**| **inheritance-disable**]

Syntax Description

| | |
|----------------------------|---|
| always | (Optional) Always performs a soft inbound clear using stored updates, even if the neighbor supports the route refresh capability. |
| inheritance-disable | (Optional) Overrides configuration for this command that may be inherited from a neighbor group or address family group. |

Command Default

Soft reconfiguration is not enabled.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

To filter or modify some of the updates received from a neighbor, you configure an inbound policy using the **route-policy (BGP)** command. Configuring soft reconfiguration inbound causes the software to store the original unmodified route beside a route that is modified or filtered. This allows a “soft clear” to be performed after the inbound policy is changed. To perform a soft clear, use the **clear bgp soft** command with the **in** keyword specified. The unmodified routes are then passed through the new policy and installed in the BGP table.

**Note**

If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.

**Note**

The **bgp auto-policy-soft-reset** is enabled by default. A soft clear is done automatically when the inbound policy configured with the **route-policy (BGP)** command is changed. This behavior can be changed by disabling the auto-policy-soft-reset using the **bgp auto-policy-soft-reset disable** command.

If the neighbor supports the route refresh capability, then the original routes are not stored because they can be retrieved from the neighbor through a route refresh request. However, if the **always** keyword is specified, the original routes are stored even when the neighbor supports the route refresh capability.

If the **soft-reconfiguration inbound** command is not configured and the neighbor does not support the route refresh capability, then an inbound soft clear is not possible. In that case, the only way to rerun the inbound policy is to use the **clear bgp ip-address** command to reset the neighbor BGP session.

**Note**

If there is an existing BGP session with a neighbor that does not support the route refresh capability, the session is terminated and a new one is initiated.

**Note**

The extra routes stored as a result of configuring this command use more memory on the router.

If you configure this command for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows inbound soft reconfiguration enabled for IP Version 4 (IPv4) unicast routes received from neighbor 10.108.1.1. The software stores all routes received in their unmodified form so that when an inbound soft clear is performed later, the stored information can then be used to generate a new set of modified routes.

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# exit
```

The following example shows inbound soft reconfiguration disabled for neighbor 10.108.1.1, preventing this feature from being automatically inherited by address family group group1:

```
RP/0/RSP0/CPU0:router(config)# router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)# af-group group1 address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# soft-reconfiguration inbound
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 10.108.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 100
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group group1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# soft-reconfiguration inbound inheritance-disable
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|---|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| bgp auto-policy-soft-reset disable , on page 54 | Disables an automatic soft reset of BGP peers when the configured inbound route policy is modified. |
| clear bgp , on page 119 | Resets a BGP connection using a soft or hard reset. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| rd , on page 227 | Applies a prefix list to filter updates received from a neighbor. |
| route-policy (BGP) , on page 249 | Applies a routing policy to updates advertised to or received from a BGP neighbor. |

speaker-id

To allocate a speaker process to a neighbor, use the **speaker-id** command in the appropriate configuration mode. To remove the speaker process from a neighbor, use the **no** form of this command.

speaker-id *id*

no speaker-id [*id*]

Syntax Description

| | |
|-----------|--|
| <i>id</i> | ID of the speaker process. Range is 1 to 15. |
|-----------|--|

Command Default

Default is 0.

Command Modes

Session group configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.0 | Removed support for this command in neighbor configuration mode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to allocate speaker process 3 to neighbor 192.168.40.24:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# speaker-id 3
```

svd platform enable

To enable selective VRF download (SVD) for Cisco ASR 9000 Series Aggregation Services Router, use the **svd platform enable** command in administrative configuration mode. To disable selective VRF download for Cisco ASR 9000 Series Aggregation Services Router, use the **no** form of this command.

svd platform enable

no svd platform enable

Syntax Description

This command has no keywords or arguments.

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

After upgrading to Cisco IOS XR Release 4.3.1 or later, the default setting for SVD is disabled and routes or labels are not selectively downloaded to the line cards. All routes are available on all line cards. This may lead to out of resource conditions, if the line card hardware scale limits are exceeded because routes were conserved by downloading selectively earlier. To resolve the out of resource condition, turn on SVD using the **svd platform enable** command.

Ensure that the total number of routes (sum of all unique routes on core facing line cards and edge facing line cards) fits on the line card hardware, before upgrading to Cisco IOS XR Release 4.3.1 or later.

To enable selective VRF download on Cisco ASR 9000 Series Aggregation Services Router, reload the chassis using the **reload location all** command after configuring the **svd platform enable** command.



Note

After enabling SVD using **svd platform enable**, do not use the **selective-vrf-download disable** command to turn off SVD.

Selective VRF download is disabled by default. Once SVD is enabled, use the **no svd platform enable** command followed by a reload of the router using the **reload location all** command to disable SVD.

Command Default

Selective VRF download is disabled.

Command Modes

Administrative configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|-----------------|------------------|
| config-services | read, write |

Examples

This example shows how to enable selective VRF download for Cisco ASR 9000 Series Aggregation Services Router:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#svd platform enable
```

table-policy

To apply a routing policy to routes being installed into the routing table, use the **table-policy** command in an appropriate configuration mode. To disable applying a routing policy when installing routes into the routing table, use the **no** form of this command.

table-policy *policy-name*

no table-policy [*policy-name*]

Syntax Description

| | |
|--------------------|--------------------------------------|
| <i>policy-name</i> | Name of the routing policy to apply. |
|--------------------|--------------------------------------|

Command Default

No policy is applied when routes are installed into the routing table.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 VRF IPv4 address family configuration
 VRF IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

Table policy provides users with the ability to drop routes from the RIB based on match criteria. This feature can be useful in certain applications and should be used with caution as it can easily create a routing 'black hole' where BGP advertises routes to neighbors that BGP does not install in its global routing table and forwarding table.

Use the **table-policy** command to modify route attributes as the routes are installed into the routing table by Border Gateway Protocol (BGP). Commonly, it is used to set the traffic index attribute.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to apply the set-traffic-index policy to IPv4 unicast routes being installed into the routing table:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1  
RP/0/RSP0/CPU0:router(config-bgp)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-bgp-af)# table-policy set-traffic-index
```

Related Commands

| Command | Description |
|---------------------------|--|
| route-policy (RPL) | Defines a route policy and enters route policy configuration mode. |

timers (BGP)

To set the timers for a specific Border Gateway Protocol (BGP) neighbor, use the **timers** command in an appropriate configuration mode. To set the timers to the default values, use the **no** form of this command.

timers *keepalive hold-time*

no timers [*keepalive hold-time*]

Syntax Description

| | |
|------------------|---|
| <i>keepalive</i> | Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535. |
| <i>hold-time</i> | Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535. |

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Use the **timers bgp** command to override the default values.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

If this command is configured for a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbours that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to change the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.40.24:

```
RP/0/RSP0/CPU0:router(config)# router bgp 109
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# timers 70 210
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| timers bgp, on page 493 | Adjusts BGP network timers for all BGP neighbors. |

timers bgp

To change the default timer values for Border Gateway Protocol (BGP) neighbors, use the **timers bgp** command in an appropriate configuration mode. To set the default timers to the default values, use the **no** form of this command.

timers bgp *keepalive hold-time*

no timers bgp [*keepalive hold-time*]

Syntax Description

| | |
|------------------|---|
| <i>keepalive</i> | Frequency (in seconds) with which the software sends keepalive messages to a neighbor. Range is 0 to 65535. |
| <i>hold-time</i> | Interval (in seconds) after not receiving a keepalive message from the neighbor that the software terminates the BGP session for the neighbor. Values are 0 or a number in the range from 3 to 65535. |

Command Default

keepalive : 60 seconds

hold-time : 180 seconds

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers bgp** command to adjust the default timer times used by all BGP neighbors. The values can be overridden on particular neighbors using the **timers** command in the neighbor configuration mode.

The timers actually used in connection with the neighbor may not be the same as those configured with this command. The actual timers are negotiated with the neighbor when establishing the session. The negotiated hold time is the minimum of the configured time and the hold time received from the neighbor. If the negotiated hold time is 0, keepalives are disabled.

The configured value for the keepalive must not exceed one-third of the negotiated hold time. If it does, a value of one-third of the negotiated hold time is used.

In cases where mechanisms such as Bi-directional Forwarding Detection (BFD), BGP fast-external-failover or Next-hop Tracking cannot be employed to detect and react to changes in the network in a faster manner, BGP Keepalive and Hold-timer values can be configured to use smaller values than the default (60 and 180 seconds respectively). When using aggressive values, consider the router's profile and scale, particularly in respect to the number of BGP neighbors that will be using sessions with the non-default timers.

Sessions using very aggressive values will be more susceptible to flap during events that cause the Route-Processor's CPU utilization levels to increase. Such events include component OIR, Route-Processor Failover, network instability, excessive churn in routing protocols etc. It is therefore recommended that the desired scale and profile of the router be tested with the non-default timer values, subjecting the router to CPU-intensive events in order to determine the timer threshold values that are appropriate for the router before configuring the values in an operational network.

The BGP Non-Stop Routing (NSR) is able to sustain sessions with more aggressive timer values than BGP Graceful Restart (GR) since in the event of a Route-Processor Failover, Graceful Restart (GR) requires the re-establishment of the TCP session over which the BGP session takes place. When using Non-Stop Routing (NSR), both the underlying TCP session and BGP session are maintained during Route-Processor failover.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure a default keepalive time of 30 seconds and a default hold time of 90 seconds:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# timers bgp 30 90
```

Related Commands

| Command | Description |
|---|--|
| timers (BGP), on page 491 | Adjusts BGP network timers for a BGP neighbor. |

transport (rpki-server)

To choose a transport mechanism for the RPKI cache-server configuration, establish and manage transport connections, and send or receive byte streams from the network, use the **transport** command in rpki-server configuration mode. To disable the transport connection, use the **no** form of this command.

transport {ssh| tcp} port *port-number*

no transport {ssh| tcp} port *port-number*

Syntax Description

| | |
|--------------------|--|
| port | Specifies to choose a port number for the RPKI cache transport. |
| <i>port-number</i> | Specifies the port number for the RPKI cache transport. Range is 1 to 65535. |

Command Default

Transport mechanism is disabled.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The transport can be set to either TCP or SSH. An SSH transport session is the recommended transport between router and RPKI cache for security reasons.

The transport method (TCP or SSH) can be configured on a per-RPKI-server basis: once server can be TCP port 980, another can be SSH port 22, for example. This can be changed by configuration. Changing the transport method will cause the cache session to flap (cleanup its existing transport related data and initialize the new transport related data).

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure SSH as the transport mechanism and to use port 1 for SSH communication:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router (config)#router bgp 100  
RP/0/RSP0/CPU0:router (config-bgp)#rpki server 172.168.35.40  
RP/0/RSP0/CPU0:router (config-bgp-rpki-cache)# transport ssh port 1
```

ttl-security

To configure a router to check the time-to-live (TTL) field in incoming IP packets for the specified external Border Gateway Protocol (eBGP) peer, use the **ttl-security** command in an appropriate configuration mode. To disable TTL verification, use the **no** form of this command.

ttl-security [**inheritance-disable**]

no ttl-security [**inheritance-disable**]

Syntax Description

| | |
|----------------------------|--|
| inheritance-disable | (Optional) Prevents the ttl-security command from being inherited from a session group or neighbor group. |
|----------------------------|--|

Command Default

TTL verification is not enabled for eBGP peers.

Command Modes

Neighbor configuration
 VRF neighbor configuration
 Neighbor group configuration
 Session group configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The disable keyword was replaced with the inheritance-disable keyword. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ttl-security** command to enable a lightweight security mechanism to protect eBGP peering sessions from CPU utilization-based and other resource exhaustion-based attacks. These types of attacks are typically brute-force Denial of Service (DoS) attacks that attempt to disable the network by flooding devices in the network with IP packets that contain forged source and destination IP addresses in the packet headers.

This command leverages existing behavior in IP packets. For a given IP packet, the TTL count of the packet always is equal to or less than the TTL count when the packet originated, a behavior that is considered impossible to circumvent. Therefore, a packet received with a TTL count equal to the maximum TTL value of 255 can be sent only by a directly adjacent peer. When the **ttl-security** command is configured for an eBGP neighbor that is directly adjacent, the router accepts only IP packets with a TTL count that is equal to the maximum TTL value.

The **ttl-security** command secures the eBGP session in the incoming direction only. In the outbound direction, it causes packets to be sent only with the maximum TTL value so that the BGP neighbor can also verify the TTL value of incoming packets. When this command is enabled, BGP establishes or maintains a session only if the TTL value in the IP packet header is equal to the maximum TTL value. If the value is less than the maximum TTL value, the packet is discarded and an Internet Control Message Protocol (ICMP) message is not generated. This behavior is designed because a response to a forged packet is not necessary.



Note The **ttl-security** command must be configured on each participating router. Failure to configure this command on both ends of the BGP session results in the session progressing as far as the OpenSent or OpenConfirm state, remaining there until the hold time expires.

The following restrictions apply to the configuration of this command:

- The **ttl-security** command should not be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The simultaneous configuration of these commands is permitted; however, the **ttl-security** command overrides the **ebgp-multihop** command.
- This command is not supported for internal BGP (iBGP) peers.
- This command is not effective against attacks from a directly adjacent peer that has been compromised.

If you configure this command for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.



Note If the **ttl-security** command is configured on a neighbor to which the router has an established connection or the router is in the process of establishing a connection, the session must be cleared using the **clear bgp** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to enable TTL security for eBGP neighbor 192.168.223.7:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.7
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65507
RP/0/RSP0/CPU0:router(config-bgp-nbr)# ttl-security
```

The following example shows how to enable TTL security for multiple eBGP neighbors using a session group:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65534
RP/0/RSP0/CPU0:router(config-bgp)# session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# ttl-security
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65501
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

```
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65502
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.223.3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65503
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group ebgp-nbrs
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

Related Commands

| Command | Description |
|---|--|
| ebgp-multihop, on page 159 | Accepts and attempts BGP connections to external peers residing on networks that are not directly connected. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |
| show lpts flows | Displays information about locally terminated packet flows, including the minimum TTL value expected. |

update limit

To set upper bound on transient memory usage for update generation, use the **update limit** command in router configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit *update-limit-MB*

no update limit

Syntax Description

| | |
|------------------------|--|
| <i>update-limit-MB</i> | Sets the update limit in megabytes (MB). Range is 16 to 2048 MB. |
|------------------------|--|

Command Default

Default update limit is 512 MB.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|--|
| Release 4.2.0 | This command was introduced and replaced the bgp write-limit command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **update limit** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the update limit as *1024* MB:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)#update limit 1024
```

Related Commands

| Command | Description |
|--|---|
| update limit address-family, on page 502 | Sets upper bound on transient memory usage for update generation for an address family. |
| update limit sub-group, on page 504 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

update limit address-family

To set upper bound on transient memory usage for update generation for an address family, use the **update limit address-family** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit address-family *update-limit-MB*

no update limit address-family

Syntax Description

| | |
|------------------------|--|
| <i>update-limit-MB</i> | Sets the update limit in megabytes (MB). Range is 4 MB to 2048 MB. |
|------------------------|--|

Command Default

Default update limit is 256 MB.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 L2VPN address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 4.2.0 | This command was introduced and replaced the bgp write-limit command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **update limit address-family** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the update limit as *512 MB* for address family VPNv6 unicast:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv6 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#update limit address-family 512
```

Related Commands

| Command | Description |
|---|---|
| update limit, on page 500 | Sets upper bound on transient memory usage for update generation. |
| update limit sub-group, on page 504 | Sets upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups. |

update limit sub-group

To set upper bound on transient memory usage for update generation for eBGP or iBGP sub-groups, use the **update limit sub-group** command in an appropriate address-family configuration mode. To return the bounds to the default value, use the **no** form of this command.

update limit sub-group {**ebgp**|**ibgp**} *update-limit-MB*

no update limit sub-group {**ebgp**|**ibgp**}

Syntax Description

| | |
|------------------------|---|
| ebgp | Specifies the update limit for eBGP sub-groups. |
| ibgp | Specifies the update limit for iBGP sub-groups. |
| <i>update-limit-MB</i> | Sets the update limit in megabytes (MB). Range is 1 MB to 512 MB. |

Command Default

Default update limit is 32 MB.

Command Modes

IPv4 address family configuration
 IPv6 address family configuration
 L2VPN address family configuration
 VPNv4 address family configuration
 VPNv6 address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 4.2.0 | This command was introduced and replaced the bgp write-limit command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **update limit sub-group** command to configure a global limit on the size of messages the software queues when updating peers. Increasing the limit can result in faster Border Gateway Protocol (BGP) convergence, but also may result in higher memory usage during convergence.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to set the update limit as *256 MB* for eBGP sub-group under address family VPNv6 unicast:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv6 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#update limit sub-group ebgp 256
```

This example shows how to set the update limit as *256 MB* for iBGP sub-group under address family IPv6 multicast:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#address-family ipv6 multicast
RP/0/RSP0/CPU0:router(config-bgp-af)#update limit sub-group ibgp 256
```

Related Commands

| Command | Description |
|---|---|
| update limit , on page 500 | Sets upper bound on transient memory usage for update generation. |
| update limit address-family , on page 502 | Sets upper bound on transient memory usage for update generation for an address family. |

update in error-handling basic disable

To disable inbound update message basic error handling for eBGP or iBGP neighbors, use the **update in error-handling basis disable** command in router configuration mode. To enable inbound update message basic error handling, use the **no** form of this command.

update in error-handling basic {ebgp|ibgp} disable

no update in error-handling basic {ebgp|ibgp} disable

Syntax Description

| | |
|-------------|---|
| ebgp | Specifies inbound update message basic error handling for eBGP neighbors. |
| ibgp | Specifies inbound update message basic error handling for iBGP neighbors. |

Command Default

Inbound update message basic error handling is enabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to disable inbound update message basic error handling for eBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#update in error-handling basic ebgp disable
```

This example shows how to disable inbound update message basic error handling for iBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
```

```
RP/0/RSP0/CPU0:router(config-bgp)#update in error-handling basic ibgp disable
```

update in error-handling extended

To enable inbound update message extended error handling for eBGP or iBGP neighbors, use the **update in error-handling extended** command in router configuration mode. To disable inbound update message error handling, use the **no** form of this command.

update in error-handling extended {ebgp|ibgp}

no update in error-handling extended {ebgp|ibgp}

Syntax Description

| | |
|-------------|--|
| ebgp | Specifies to enable inbound update message extended error handling for eBGP neighbors. |
| ibgp | specifies to enable inbound update message extended error handling for iBGP neighbors. |

Command Default

Inbound update message extended error handling is disabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to enable inbound update message extended error handling for eBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#update in error-handling extended ebgp
```

This example shows how to enable inbound update message extended error handling for iBGP neighbors:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#router bgp 100  
RP/0/RSP0/CPU0:router(config-bgp)#update in error-handling extended ibgp
```

update out logging

To enable logging of update generation events, use the **update out logging** command in router configuration mode. To disable the logging of update generation events, use the **no** form of this command.

update out logging

no update out logging

Syntax Description This command has no keywords or arguments.

Command Default Update generation event logging is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | bgp | read, write |

Examples This example shows how to enable logging of update generation events:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#update out logging
```

update-source

To allow internal Border Gateway Protocol (iBGP) sessions to use the primary IP address from a particular interface as the local address when forming an iBGP session with a neighbor, use the **update-source** command in an appropriate configuration mode. To set the chosen local IP address to the nearest interface to the neighbor, use the **no** form of this command.

update-source *type interface-path-id*

no update-source [*type interface-path-id*]

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

Best local address

Command Modes

Neighbor configuration
VRF neighbor configuration
Neighbor group configuration
Session group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **update-source** command is commonly used with the loopback interface feature for iBGP sessions. The loopback interface is defined, and the interface address is used as the endpoint for a BGP session through the **update-source** command. This mechanism allows a BGP session to remain up even if the outbound interface goes down, provided there is another route to the neighbor.

If this command is configured for a neighbor group or session group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to configure this router to use the IP address from the Loopback0 interface when trying to open a session with neighbor 172.20.16.6:

```
RP/0/RSP0/CPU0:router(config)# router bgp 110
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 110
RP/0/RSP0/CPU0:router(config-bgp-nbr)# update-source Loopback0
```

Related Commands

| Command | Description |
|---|--|
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |

update wait-install

To configure BGP to wait for feedback from RIB indicating that the routes that BGP installed in RIB have been installed in FIB, before BGP send out updates to neighbors, use the **update wait-install** command in an appropriate configuration mode.

update wait-install

no update wait-install

This command has no keywords or arguments.

Command Default

The update wait-install configuration is disabled.

Command Modes

Router IPv4 address family

Router VPNv4 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to enable the **update wait-install** configuration under VPNv4 unicast address family:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 65500
RP/0/RSP0/CPU0:router(config-bgp)#address-family vpnv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-af)#update wait-install
```

USE

To inherit configuration from a neighbor group, session group, or address family group, use the **use** command in an appropriate configuration mode. To discontinue inheritance from a group, use the **no** form of this command.

```
use {af-group group-name| neighbor-group group-name| session-group group-name}
no use {af-group [group-name]| neighbor-group [group-name]| session-group [group-name]}
```

Syntax Description

| | |
|-----------------------|--|
| af-group | Specifies an address family group. |
| <i>group-name</i> | Name of the neighbor group, session group, or address family group from which you want to inherit configuration. |
| neighbor-group | Specifies a neighbor group. |
| session-group | Specifies a session group. |

Command Default

Inheritance of group characteristics does not occur.

Command Modes

For **use af-group** version:
 Address family group configuration
 Neighbor address family configuration
 Neighbor group address family configuration

For **use neighbor-group** version:
 Neighbor group configuration
 Neighbor configuration
 VRF neighbor configuration

For **use session-group** version:
 Neighbor group configuration
 Neighbor configuration
 VRF neighbor configuration
 Session-group configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **use** command configures inheritance of configuration from an address family group, neighbor group, or session group, which means that any configuration for the group also takes effect for the user of the group.

The configuration inherited depends on the type of group that is specified. The group types are described in the following sections:

Address Family Group

An address family group can specify a configuration for only a single address family. The address family specified when the address family group was defined (through the **af-group** command) must match the address family from which the group is used.

Neighbor Group

A neighbor group (like a neighbor) can have address family-independent configuration and address family-specific configuration. All of these configurations could be inherited.

Session Group

A session group can have only address family-independent configuration and thus only address family-independent configuration is inherited from it.

The following rules govern inheritance to resolve possible conflicting configuration:

- 1 If a command is configured directly on the neighbor that is using group configuration, the command overrides the value that would be normally inherited from the group.
- 2 If the neighbor is configured to use a session group (for address family-independent configuration) or an address family group (for address family-specific configuration) and the command is configured for the session group or address family group, that configuration is used.
- 3 The neighbor group configuration is used:
 - If the command is not configured directly on the neighbor and the neighbor is not using a session group (for address family-independent configuration) or an af-group (for address family-specific configuration).
 - The neighbor is using a neighbor group and the command is configured on the neighbor group.

Typically, all configuration for a neighbor group is inherited, but some characteristics may be masked by a session group or address family group. For an example of this configuration, see the “Examples” section.

If the neighbor is using both a session group and a neighbor group and a specific command is configured for the neighbor group but not for the session group, then the configuration for the neighbor group does not take effect. The session group “hides” all address family-independent configuration on the neighbor group and prevents it from being inherited. Similarly, the use of an address family group hides any address family-specific configuration that may otherwise be inherited from a neighbor group for that address family.

In addition to neighbors using groups, it is possible to build a hierarchy by having groups use other groups. The following hierarchical groups are permitted:

- Session groups may use other session groups.
- Address family groups may use other address family groups.

- Neighbor groups may use other neighbor groups.
- Neighbor groups may use session groups and address family groups.

**Note**

Within the Cisco IOS XR system configuration architecture, do not combine the **remote-as** command and the **no use neighbor-group** command in the same commit, or the **remote-as** command and the **no use session-group** command in the same commit.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to define a session group session1 and configure neighbor 172.168.40.24 to use session1. As a result, the session1 configuration takes effect on the neighbor also.

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group session1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

The following example is similar to the previous example, but in this case the **timers** command on the session group does not take effect on the neighbor because it is overridden by a **timers** command directly configured for the neighbor.

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group session1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# timers 30 90
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group session1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# timers 60 180
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit
```

The following example shows an address family group, family1, for IPv4 multicast and a neighbor group, neighbor1, that have IPv4 unicast and IPv4 multicast enabled. In this case, the neighbor inherits IPv4 unicast (and address family-independent) configuration from the neighbor group, but inherits IPv4 multicast configuration from the address family group. In this example, the neighbor group also has a remote autonomous system configured, so there is no need to configure a remote autonomous system for the neighbor because it inherits the remote autonomous system from the neighbor group:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# af-group family1 address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# route-policy mcast-in in
RP/0/RSP0/CPU0:router(config-bgp-afgrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group neighbor1
```

```

RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 in
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# route-policy policy1 out
RP/0/RSP0/CPU0:router(config-bgp-nbrgrp-af)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group neighbor1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 multicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# use af-group family1
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# exit

```

In the previous example, the neighbor uses the policy1 route policy for inbound and outbound IPv4 unicast routes, but uses the mcast-in route policy for inbound IPv4 multicast routes and no policy for outbound IPv4 multicast routes.

The following example shows a neighbor inheriting configuration from a session group that likewise inherits configuration from another session group. The configuration from both session groups take effect on the neighbor:

```

RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# session-group session1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 40
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# session-group session2
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# use session-group session1
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# update-source Loopback0
RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.168.40.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group session2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# exit

```

Related Commands

| Command | Description |
|---|---|
| af-group , on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| session-group , on page 264 | Creates a session group and enters session group configuration mode. |
| neighbor-group , on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| remote-as (BGP) , on page 239 | Creates a BGP neighbor and begins the exchange of routing information. |
| show bgp af-group , on page 301 | Displays information about BGP configuration for address family groups. |
| show bgp neighbor-group , on page 349 | Displays information about the BGP configuration for neighbor groups. |
| show bgp neighbors , on page 353 | Displays information about BGP neighbors. |

| Command | Description |
|--|--|
| show bgp session-group , on page 436 | Displays information about the BGP configuration for session groups. |

username (rpki-server)

To specify a SSH **username** for the RPKI cache-server, use the **username** command in rpki-server configuration mode. To remove the username, use the **no** form of this command.

username *user-name*

no username *user-name*

Syntax Description

| | |
|------------------|---|
| <i>user-name</i> | Enters a username to be used for the SSH transport mechanism. |
|------------------|---|

Command Default

Username is not configured.

Command Modes

RPKI server configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The username configuration applies only if the SSH transport mechanism is active.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |

Examples

This example shows how to configure a username (*rpki-user*) for the RPKI cache-server SSH transport mechanism:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router bgp 100
RP/0/RSP0/CPU0:router(config-bgp)#rpki server 172.168.35.40
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)# transport ssh port 1
RP/0/RSP0/CPU0:router(config-bgp-rpki-cache)#username rpki-user
```

vrf (BGP)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove the VRF instance from the configuration file and restore the system to its default condition, use the **no** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

| Syntax Description | |
|--------------------|---|
| <i>vrf-name</i> | Name of the VRF instance. The following names cannot be used: all, default, and global. |

| Command Default | No default behavior or values |
|-----------------|-------------------------------|
|-----------------|-------------------------------|

| Command Modes | Router configuration |
|---------------|----------------------|
|---------------|----------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

Examples

The following example shows how to configure a VRF instance and enter VRF configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-bgp-vrf)#
```

weight

To assign a weight to routes received from a neighbor, use the **weight** command in an appropriate configuration mode. To remove the **weight** command from the configuration file and restore the system to its default condition in which the software assigns the default weight to routes, use the **no** form of this command.

weight *weight-value*

no weight [*weight-value*]

Syntax Description

weight-value

Weight to assign. Range is 0 to 65535.

Command Default

Routes learned through another Border Gateway Protocol (BGP) peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768.

Command Modes

IPv4 address family group configuration
 IPv6 address family group configuration
 VPNv4 address family group configuration
 IPv4 neighbor address family configuration
 VPNv4 neighbor address family configuration
 VRF IPv4 neighbor address family configuration
 IPv4 neighbor group address family configuration
 IPv6 neighbor group address family configuration
 VPNv4 neighbor group address family configuration
 VRF IPv6 neighbor address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The weight of a route is a Cisco-specific attribute. It is used in the best-path selection process (as the strongest tie-breaker). See the *Implementing BGP on Cisco ASR 9000 Series Router* module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide* for information on best path. If there are two BGP routes with the same network layer reachability information (NLRI), the route with the higher weight is always chosen no matter what the value of other BGP attributes. Weight only has significance on the local

router. Weight is assigned locally to the router, is a value that only makes sense to the specific router, is not propagated or carried through any route updates, and never is sent between BGP peers (even within the same AS).

**Note**

If an address family group, neighbor group, or session group is configured, the configuration inside these configuration groups will not be effective unless it is applied directly or indirectly to one or more neighbors.

The weight assigned to individual routes can be further manipulated in the inbound route policy of a neighbor using the **set weight** command. The **set weight** command sets the weight directly. If you have particular neighbors that you want to prefer for most of your outbound traffic, you can assign a higher weight to all routes learned from that neighbor.

The weight assigned to individual routes may be modified by using an inbound routing policy.

**Note**

For weight changes to take effect, you may need to use the [clear bgp soft, on page 137](#) command.

If this command configures a neighbor group or neighbor address family group, all neighbors using the group inherit the configuration. Values of commands configured specifically for a neighbor override inherited values.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |

Examples

The following example shows how to assign a weight of 50 to all IP Version 4 (IPv4) unicast routes learned through 172.20.16.6:

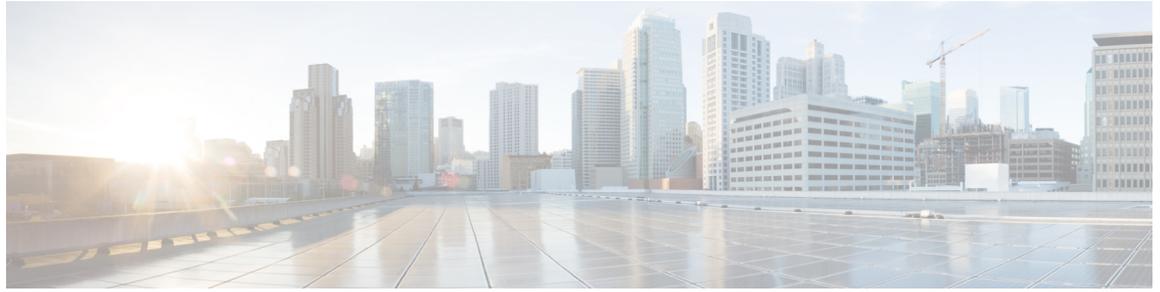
```
RP/0/RSP0/CPU0:router(config)# router bgp 1
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 172.20.16.6
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# weight 50
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# exit
```

Related Commands

| Command | Description |
|---|---|
| af-group, on page 26 | Creates an address family group for BGP neighbors and enters address family group configuration mode. |
| clear bgp, on page 119 | Resets a group of BGP neighbors. |
| neighbor-group, on page 194 | Creates a neighbor group and enters neighbor group configuration mode. |
| session-group, on page 264 | Creates a session group and enters session group configuration mode. |

| Command | Description |
|------------|---------------------------------|
| set weight | Sets the weight for BGP routes. |

 weight



BFD Commands

This module provides command line interface (CLI) commands for configuring Bidirectional Forwarding Detection (BFD) on the Cisco ASR 9000 Series Router.

- [address-family ipv4 unicast \(BFD\)](#), page 527
- [bfd](#), page 529
- [bfd address-family ipv4 destination](#), page 531
- [bfd address-family ipv4 fast-detect](#), page 533
- [bfd address-family ipv4 minimum-interval](#), page 534
- [bfd address-family ipv4 multiplier](#), page 537
- [bfd address-family ipv4 timers](#), page 540
- [bundle coexistence bob-blb](#), page 542
- [bfd dampening](#), page 544
- [bfd fast-detect](#), page 546
- [bfd minimum-interval](#), page 549
- [bfd multipath include location](#), page 553
- [bfd multiplier](#), page 555
- [clear bfd counters](#), page 558
- [dampening \(BFD\)](#), page 560
- [echo disable](#), page 563
- [echo ipv4 source](#), page 565
- [echo latency detect](#), page 567
- [echo startup validate](#), page 569
- [interface \(BFD\)](#), page 571
- [ipv6 checksum](#), page 574
- [multihop ttl-drop-threshold](#), page 576

- [show bfd](#), page 578
- [show bfd client](#), page 581
- [show bfd counters](#), page 583
- [show bfd mib session](#), page 586
- [show bfd multipath](#), page 589
- [show bfd neighbor](#), page 591
- [show bfd session](#), page 593
- [show bfd summary](#), page 601

address-family ipv4 unicast (BFD)

To enable Bidirectional Forwarding Detection (BFD) fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address, use the **address-family ipv4 unicast** command in static route configuration mode. To return the router to the default setting, use the **no** form of this command.

address-family ipv4 unicast *address nexthop* **bfd fast-detect** [**minimum interval** *interval*] [**multiplier** *multiplier*]

no address-family ipv4 unicast *address nexthop* **bfd fast-detect** [**minimum interval** *interval*] [**multiplier** *multiplier*]

Syntax Description

| | |
|---|---|
| <i>address</i> | Specifies the IPv4 unicast destination address and prefix on which to enable BFD fast-detection. |
| <i>nexthop</i> | Specifies the next-hop address on which to enable BFD fast-detection. |
| bfd fast-detect | Enables BFD fast-detection on the specified IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| minimum interval <i>interval</i> | (Optional) Ensures that the next hop is assigned with the same hello interval. Replace <i>interval</i> with a number that specifies the interval in milliseconds. Range is from 10 through 10000. |
| multiplier <i>multiplier</i> | (Optional) Ensures that the next hop is assigned with the same detect multiplier. Replace <i>multiplier</i> with a number that specifies the detect multiplier. Range is from 1 through 10. |

Command Default

interval: 100
multiplier: 3

Command Modes

Static route configuration mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the multiplier is changed using the **bfd multiplier** command, the new parameter is used to update all existing BFD sessions for the protocol (BGP, IS-IS, MPLS-TE, or OSPF).

Task ID

| Task ID | Operations |
|---------|-------------|
| static | read, write |

Examples

The following example shows how to enable BFD on a static route. In this example, BFD sessions are established with the next-hop 3.3.3.3 when it becomes reachable.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router (config-static)# address-family ipv4 unicast 2.2.2.0/24 3.3.3.3 bfd fast-detection
```

Related Commands

| Command | Description |
|--|---|
| bfd fast-detect, on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

bfd

To enter Bidirectional Forwarding Detection (BFD) configuration mode, use the **bfd** command in global configuration mode. To exit BFD configuration mode and return to global configuration mode, use the **no** form of this command.

bfd

no bfd

Syntax Description This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you issue the **bfd** command in global configuration mode, the CLI prompt changes to “config-bfd,” indicating that you have entered BFD configuration mode. In the following sample output, the question mark (?) online help function displays all the commands available under BFD configuration mode:

```
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# ?

commit      Commit the configuration changes to running
describe    Describe a command without taking real actions
do          Run an exec command
echo        Configure BFD echo parameters
exit        Exit from this submode
interface    Configure BFD on an interface
no          Negate a command or set its defaults
root        Exit to the global
            configuration mode
show        Show contents of configuration
```

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to enter BFD configuration mode:

```
RP/0/RSP0/CPU0:router # configure
RP/0/RSP0/CPU0:router (config) # bfd
RP/0/RSP0/CPU0:router (config-bfd) #
```

Related Commands

| Command | Description |
|--|---|
| echo disable, on page 563 | Disables echo mode on a router or on an individual interface or bundle. |
| interface (BFD), on page 571 | Enters BFD interface configuration mode, where you can disable echo mode on an interface. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

bfd address-family ipv4 destination

To specify the destination address for BFD sessions on bundle member links, use the **bfd address-family ipv4 destination** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 destination *ip-address*

no bfd address-family ipv4 destination *ip-address*

Syntax Description

| | |
|-------------------|---|
| <i>ip-address</i> | 32-bit IPv4 address in dotted-decimal format (A.B.C.D). |
|-------------------|---|

Command Default

No destination IPv4 address is configured.

Command Modes

Interface configuration (config-if)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is supported on bundle interfaces only.

Task ID

| Task ID | Operations |
|---------|-------------|
| bundle | read, write |

Examples

The following example specifies the IPv4 address of 10.20.20.1 as the destination address for the BFD session on an Ethernet bundle interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 destination 10.20.20.1
```

Related Commands

| Command | Description |
|--|---|
| bfd address-family ipv4 fast-detect, on page 533 | Enables IPv4 BFD sessions on bundle member links. |

bfd address-family ipv4 fast-detect

To enable IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 fast-detect** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 fast-detect

no bfd address-family ipv4 fast-detect

Syntax Description This command has no keywords or arguments.

Command Default BFD sessions are disabled.

Command Modes Interface configuration (config-if)

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is supported on bundle interfaces only.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bundle | read, write |

Examples The following example enables IPv4 BFD sessions on member links of an Ethernet bundle:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 fast-detect
```

| Related Commands | Command | Description |
|------------------|---|--|
| | bfd address-family ipv4 destination , on page 531 | Specifies the destination address for BFD sessions on bundle member links. |

bfd address-family ipv4 minimum-interval

To specify the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 minimum-interval** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 minimum-interval *milliseconds*

no bfd address-family ipv4 minimum-interval [*milliseconds*]

Syntax Description

milliseconds Shortest interval between sending BFD control packets to a neighbor. The range is 15 to 30000 milliseconds.

Note Although the command allows you to configure a minimum of 15 ms, the supported minimum is 50 ms.

Command Default

The default is 150 ms.

Command Modes

Interface configuration (config-if)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is supported on bundle interfaces only.

The BFD minimum interval is used with a configurable multiplier (**bfd address-family ipv4 multiplier** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of I and a multiplier of M , the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of I —Minimum period between sending of BFD control packets.
- Value of $I \times M$
 - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
 - Minimum period between sending of BFD echo packets.

- Value of $(I \times M) \times M$ —BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

When used with bundled VLANs, the following restrictions apply:

- The command specifies control packet intervals only because echo packets are not supported.
- The minimum interval is 250 ms.

The **bfd address-family ipv4 minimum-interval** command in bundle interface configuration overrides the minimum intervals specified by the **bfd minimum-interval** command in other areas of BFD configuration.



Note

When multiple applications share the same BFD session, the application with the most aggressive timer is used locally. Then, the result is negotiated with the peer router.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
 - The maximum rate for BFD sessions per linecard is 9600 pps.
- The maximum number of all BFD sessions on the router is 1024.
- The maximum number of all BFD sessions on the router is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:
Asynchronous rate per bundle member = $(1000 / \text{Min-interval})$

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).
Echo interval = $(\text{Min-interval} \times \text{Multiplier})$
- Calculate the overall rate supported for all members on the bundle:
Ethernet bundle rate = $(1000 / \text{Echo interval}) \times 64$
- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:
Total bundle rate = Ethernet bundle rate + $(\text{Base asynchronous rate} \times \text{Number of links})$

Task ID

| Task ID | Operations |
|---------|-------------|
| bundle | read, write |

Examples

The following example specifies that control packets will be sent at a minimum interval of 200 ms for IPv4 BFD sessions on member links of an Ethernet bundle:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
```

Related Commands

| Command | Description |
|--|---|
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd address-family ipv4 multiplier , on page 537 | Specifies a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links. |

bfd address-family ipv4 multiplier

To specify a number that is used as a multiplier with the minimum interval to determine BFD control and echo packet failure detection times and echo packet transmission intervals for IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 multiplier** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 multiplier *multiplier*

no bfd address-family ipv4 multiplier [*multiplier*]

Syntax Description

| | |
|-------------------|--|
| <i>multiplier</i> | Number from 2 to 50. |
| Note | Although the command allows you to configure a minimum of 2, the supported minimum is 3. |

Command Default

The default multiplier is 3.

Command Modes

Interface configuration (config-if)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is supported on bundle interfaces only.

The BFD multiplier is used with a configurable minimum interval (**bfd address-family ipv4 minimum-interval** command) to determine the intervals and failure detection times for both control and echo packets in asynchronous mode on bundle member links.

For example, with a session interval of I and a multiplier of M , the following packet intervals and failure detection times apply for BFD asynchronous mode:

- Value of I —Minimum period between sending of BFD control packets.
- Value of $I \times M$
 - BFD control packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD control packet before the session is declared down.
 - Minimum period between sending of BFD echo packets.



Note The maximum echo packet interval for BFD on bundle member links is the minimum of either 30 seconds or the asynchronous control packet failure detection time.

- Value of $(I \times M) \times M$ —BFD echo packet failure detection time. This is the maximum amount of time that can elapse without receipt of a BFD echo packet before the session is declared down.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
 - The maximum rate for BFD sessions per linecard is 9600 pps.
- The maximum number of all BFD sessions per linecard is 1024.
- The maximum number of all BFD sessions per linecard is 1440.

To calculate the rate for BFD sessions on bundle members running in asynchronous mode without echo:

- Divide 1000 by the value of the minimum interval (as specified by the **bfd address-family ipv4 minimum-interval** command). This is also the base rate used per member session with echo:
Asynchronous rate per bundle member = $(1000 / \text{Min-interval})$

To calculate the rate for BFD sessions on bundle members running in asynchronous mode with echo:

- Determine the echo interval, which is the value of the minimum interval (specified by the **bfd address-family ipv4 minimum-interval** command) multiplied by the multiplier value (specified by the **bfd address-family ipv4 multiplier** command).
Echo interval = $(\text{Min-interval} \times \text{Multiplier})$
- Calculate the overall rate supported for all members on the bundle:
Ethernet bundle rate = $(1000 / \text{Echo interval}) \times 64$
- Add the asynchronous base rate per bundle member to find the total rate for all bundle links:
Total bundle rate = Ethernet bundle rate + $(\text{Base asynchronous rate} \times \text{Number of links})$

Task ID

| Task ID | Operations |
|---------|-------------|
| bundle | read, write |

Task ID

Examples

The following example specifies the following packet intervals and failure detection times for IPv4 BFD sessions on member links with asynchronous echo mode on an Ethernet bundle:

- 200 ms control packet interval
- 600 ms control packet failure detection interval

- 600 ms echo packet interval
- 1800 ms echo packet failure detection interval

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 minimum-interval 200
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 multiplier 3
```

Related Commands

| Command | Description |
|--|---|
| bfd address-family ipv4 minimum-interval , on page 534 | Specifies the minimum interval for asynchronous mode control packets on IPv4 BFD sessions on bundle member links. |
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |

bfd address-family ipv4 timers

To configure timers to allow for delays in receipt of BFD state change notifications (SCNs) from peers before declaring a link bundle BFD session down for IPv4 BFD sessions on bundle member links, use the **bfd address-family ipv4 timers** command in interface configuration mode. To return to the default, use the **no** form of this command.

bfd address-family ipv4 timers [**start**| **nbr-unconfig**] *seconds*

no bfd address-family ipv4 timers [**start**| **nbr-unconfig**] *seconds*

Syntax Description

| | |
|------------------------------------|---|
| start <i>seconds</i> | Number of seconds after startup of a BFD member link session to wait for the expected notification from the BFD peer to be received, so that the session can be declared up. If the SCN is not received after that period of time, the BFD session is declared down. The range is 60 to 3600. Note In Cisco IOS XR Releases 4.0 and 4.0.1, the available minimum is 30, but is not recommended. |
| nbr-unconfig <i>seconds</i> | Number of seconds to wait after receipt of notification that the BFD configuration has been removed by a BFD neighbor, so that any configuration inconsistency between the BFD peers can be fixed. If the BFD configuration issue is not resolved before the specified timer is reached, the BFD session is declared down. The range is 60 to 3600. Note In Cisco IOS XR Releases 4.0 and 4.0.1, the available minimum is 30, but is not recommended. |

Command Default

No timers are configured.

Command Modes

Interface configuration (config-if)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command is supported on bundle interfaces only.

Task ID

| Task ID | Operations |
|---------|-------------|
| bundle | read, write |

Examples

The following example configures a timer for members of the specified Ethernet bundle that allows up to 1 minute (60 seconds) after startup of a BFD member link session to wait for receipt of the expected notification from the BFD peer to declare the session up:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1  
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 timers start 60
```

The following example configures a timer for members of the specified Ethernet bundle that allows up to 1 minute (60 seconds) to wait after receipt of notification that the BFD configuration has been removed by a BFD neighbor, before declaring a BFD session down:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 1  
RP/0/RSP0/CPU0:router(config-if)# bfd address-family ipv4 timers nbr-unconfig 60
```

bundle coexistence bob-blb

To configure a coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB), use the **bundle coexistence bob-blb** command in BFD configuration mode. To disable this feature use the **no** form of this command.

bundle coexistence bob-blb {inherit| logical}

no bundle coexistence bob-blb {inherit| logical}

Syntax Description

| | |
|----------------|---|
| inherit | Specifies inheritance as the coexistence mechanism. |
| logical | Specifies to use BFD logical bundle natively for coexistence. |

Command Default

No coexistence configuration is enabled. The BLB configuration behaves the same as the BVLAN configuration.

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The default behavior for BLB without any coexistence configuration behaves exactly as the BVLAN configuration.

When the "inherited" coexistence mode is configured, the BLB session always uses the inherited mode. It always creates a virtual session, and never a BFD session with real packets. This mode prevents the automatic conversion between BLB and the inherited mode, which could happen in default mode, depending on the sequence in which the configuration is applied. The "inherited" session is in "Down" state when BoB is not enabled.

The "logical" mode always creates BFD session with real packets.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |

| Task ID | Operation |
|---------|-------------|
| isis | read, write |
| mpls-te | read, write |

Examples

This example shows how to configure the coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "inherited":

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#bfd
RP/0/RSP0/CPU0:router(config-bfd)#bundle coexistence bob-blb inherit
```

This example shows how to configure the coexistence mechanism between BFD over Bundle (BoB) and BFD over Logical Bundle (BLB) as being "logical":

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#bfd
RP/0/RSP0/CPU0:router(config-bfd)#bundle coexistence bob-blb logical
```

bfd dampening

To configure a device to dampen a flapping Bidirectional Forwarding Detection (BFD) session, use the **bfd dampening** command in global configuration mode. To disable dampening, use the **no** form of this command.

bfd dampening bundle-member {**initial-wait** | **l3-only-mode** | **maximum-wait** | **secondary-wait**} **initial-wait** **maximum-wait** **secondary-wait** *milliseconds*

no bfd dampening

Syntax Description

| | |
|-----------------------|---|
| bundle-member | Specifies initial, maximum, or secondary delays in milliseconds for BFD session startup on BFD bundle members. |
| initial-wait | Specifies the initial delay in milliseconds before starting a BFD session. For bundle members, the default is 16000. For non-bundle interfaces, the default is 2000. |
| maximum-wait | Specifies the maximum delay in milliseconds before starting a BFD session. For bundle members, the default is 600000. For non-bundle interfaces, the default is 12000. Note The maximum delay must be greater than the initial delay. |
| secondary-wait | Specifies a secondary delay in milliseconds before starting a BFD session. For bundle members, the default is 20000. For non-bundle interfaces, the default is 5000. |
| l3-only-mode | Applies dampening, if the failure is Layer 3 (L3)-specific. |
| <i>milliseconds</i> | For bundle members, the range is 1-518400000. For non-bundle interfaces, the range is 1-3600000. |

Command Default

A flapping BFD session on a device is not dampened.

Command Modes

Global configuration

Command History

| Release | Modification |
|-------------|------------------------------|
| Release 5.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

This example shows how to configure BFD dampening by specifying an initial and maximum delay for BFD session startup on BFD bundle members:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd dampening bundle-member initial-wait 1000
RP/0/RSP0/CPU0:router(config)# bfd dampening bundle-member maximum-wait 3000
```

This example shows how to configure BFD dampening on a non-bundle interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd dampening initial-wait 1000
RP/0/RSP0/CPU0:router(config)# bfd dampening maximum-wait 3000
```

bfd fast-detect

To enable Bidirectional Forwarding Detection (BFD) to detect failures in the path between adjacent forwarding engines, use the **bfd fast-detect** command in the appropriate configuration mode. To return the software to the default state in which BFD is not enabled, use the **no** form of this command.

bfd fast-detect [**disable**] **ipv4**

no bfd fast-detect

Syntax Description

| | |
|----------------|--|
| disable | Disables the detection of failures in the path between adjacent forwarding engines for a specified entity, such as a BGP neighbor or OSPF interface. Note The disable keyword is available in the following modes only: BGP configuration, OSPF area configuration, OSPF area interface configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration . |
| ipv4 | Enables Intermediate System-to-Intermediate System (IS-IS) BFD detection of failures in the path between adjacent forwarding engines. Note The ipv4 keyword is available in IS-IS router configuration mode only. |

Command Default

BFD detection of failures in the path between adjacent forwarding engines is disabled.

Command Modes

Neighbor configuration
 Session group configuration
 Neighbor group configuration
 Interface configuration
 Interface configuration
 Router configuration
 Area configuration
 Area interface configuration
 Router configuration
 Area configuration
 Area interface configuration
 Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|---|
| Release 4.0.0 | The bfd fast-detect command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

BFD can support multihop for internal and external BGP peers.

Use the **bfd fast-detect** command to provide protocol- and media-independent, short-duration failure detection of the path between adjacent forwarding engines, including the interfaces and data links.

BFD must be configured on directly connected neighbors for a BFD session to be established between the neighbors.

When MPLS-TE tunnels are protected by backup tunnels, BFD failure triggers fast reroute on affected tunnels.

In OSPF and OSPFv3 environments, the setting of the **bfd fast-detect** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you enable BFD in area interface configuration mode, it is enabled on the specified interface only.
- If you enable BFD in area configuration mode, it is enabled on all interfaces in the specified area.
- If you enable BFD in router configuration mode, it is enabled on all areas and all associated interfaces in the specified routing process.

The **disable** keyword is available in the following modes: BGP configuration, OSPF area configuration, OSPF area interface configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration. In OSPF and OSPFv3 environments, the **disable** option enables you to override the inheritance rules described previously. For example, if you enable BFD in an OSPF area, BFD is enabled on all interfaces in that area. If you do not want BFD running on one of the interfaces in that area, you must specify the **bfd fast-detect disable** command for that interface only.

To disable BFD or return the software to the default state in which BFD is not enabled in IS-IS router configuration mode and MPLS-TE configuration mode, you must enter the **no bfd fast-detect** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| isis | read, write |
| mpls-te | read, write |
| ospf | read, write |

| Task ID | Operations |
|-----------|-------------|
| multicast | read, write |

Examples

The following example shows how to configure BFD on a BGP router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.70.24
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 2
RP/0/RSP0/CPU0:router(config-bgp-nbr)# bfd fast-detect
```

The following example shows the configuration of an OSPFv3 routing process named `san_jose`. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, BFD is enabled at the area level, which means that by virtue of the inheritance rules, BFD is enabled on all interfaces within the area except those on which BFD is explicitly disabled. Given this rule, BFD is enabled on Gigabit Ethernet interfaces 1/0/0/0 and 2/0/0/0 and is disabled on interface 3/0/0/0.

In area 1, BFD is enabled on Gigabit Ethernet interface 5/0/0/0 only. Because BFD is not enabled at the area level nor explicitly enabled on interface 4/0/0/0, it is disabled on this interface.

```
router ospfv3 san_jose
  area 0
    bfd fast-detect
    ...
    int gige 1/0/0/0
    !
    int gige 2/0/0/0
    ...
int gige 3/0/0/0
    bfd fast-detect disable
!
!
  area 1
    int gige 4/0/0/0
    !
    int gige 5/0/0/0
    bfd fast-detect
    !
!
```

Related Commands

| Command | Description |
|--|--|
| address-family ipv4 unicast (BFD), on page 527 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| bfd minimum-interval, on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 555 | Sets the BFD multiplier. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

bfd minimum-interval

To specify the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope, use the **bfd minimum-interval** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd minimum-interval *milliseconds*

no bfd minimum-interval [*milliseconds*]

Syntax Description

| | |
|---------------------|--|
| <i>milliseconds</i> | Interval between sending BFD hello packets to the neighbor. The range is 15 to 30000 milliseconds. For MPLS-TE, the range is 15 to 200 milliseconds. |
|---------------------|--|

Command Default

BGP *interval*: 50 milliseconds

IS-IS *interval*: 150 milliseconds

OSPF and OSPFv3 *interval*: 150 milliseconds

MPLS-TE *interval*: 15 milliseconds

PIM *interval*: 150 milliseconds

Command Modes

Router configuration

Interface configuration

MPLS TE configuration

Router configuration

Area configuration

Area interface configuration

Router configuration

Area configuration

Interface configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | The bfd minimum-interval command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In OSPF and OSPFv3 environments, the setting of the **bfd minimum-interval** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you configure the minimum interval in area interface configuration mode, the updated interval affects the BFD sessions on the specified interface only.
- If you configure the minimum interval in area configuration mode, the updated interval affects the BFD sessions on all interfaces in the specified area.
- If you configure the minimum interval in router configuration mode, the updated interval affects the BFD sessions in all areas and all associated interfaces in the specified routing process.

If desired, you can override these inheritance rules by explicitly configuring the **bfd minimum-interval** command for a specific area interface or area.

**Note**

When multiple applications share the same BFD session, the application with the most aggressive timer wins locally. Then, the result is negotiated with the peer router.

Keep the following router-specific rules in mind when configuring the minimum BFD interval:

- The maximum rate in packets-per-second (pps) for BFD sessions is linecard-dependent. If you have multiple linecards supporting BFD, then the maximum rate for BFD sessions per system is the supported linecard rate multiplied by the number of linecards.
 - The maximum rate for BFD sessions per linecard is 9600 pps.
- If a session is running in asynchronous mode without echo, then PPS used for this session is (1000 / asynchronous interval in milliseconds).
- If a session is running in asynchronous mode with echo, then PPS used for this session is (1000 / echo interval in milliseconds).
This is calculated as: 1000 / value of the **bfd minimum-interval** command.

**Note**

The rate for BFD sessions on bundle member links is calculated differently. For more information, see the **bfd address-family ipv4 minimum-interval** command.

- The maximum number of all BFD sessions per linecard is 1024.
- The maximum number of all BFD sessions per linecard is 1440.
- When asynchronous mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1024 sessions, the failure detection interval must be greater than or equal to 150 milliseconds.

- When asynchronous mode is available, the minimum interval must be greater than or equal to 250 milliseconds, with a multiplier of 3 for up to 100 sessions per line card
- When asynchronous mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1440 sessions, the failure detection interval must be greater than or equal to 150 milliseconds.
- When echo mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1024 sessions, the failure detection interval must be less than or equal to 150 milliseconds.
- When echo mode is available, the minimum interval must be 50 milliseconds with a multiplier of 3.
- When echo mode is available, the minimum interval must be greater than or equal to 15 milliseconds for up to 100 sessions on the line card. If you are running the maximum of 1440 sessions, the failure detection interval must be less than or equal to 150 milliseconds.

Task ID

| Task ID | Operations |
|-----------|-------------|
| bgp | read, write |
| isis | read, write |
| mpls-te | read, write |
| ospf | read, write |
| multicast | read, write |

Examples

The following example shows how to set the BFD minimum interval for a BGP routing process:

```
RP/0/RSP0/CPU0:router(config)# router bgp 6500
RP/0/RSP0/CPU0:router(config-bgp)# bfd minimum-interval 275
```

The following example shows the configuration of an OSPFv3 routing process named san_jose. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, the minimum interval is set to 200 at the area level, which means that by virtue of the inheritance rules, the same value is set on all interfaces within the area except those on which a different value is explicitly configured. Given this rule, Gigabit Ethernet interface 1/0/0/0 uses the interval of 200, which is inherited from the area, while interface 2/0/0/0 uses the explicitly configured value of 300.

In area 1, the minimum interval is not configured at the area or interface levels, which means that interfaces 3/0/0/0 and 4/0/0/0 use the default interval of 150.

```
router ospfv3 san_jose
bfd fast-detect
  area 0
bfd minimum-interval 200
int gige 1/0/0/0
!
int gige 2/0/0/0
bfd minimum-interval 300
!
!
```

```

    area 1
  int gige 3/0/0/0
    !
  int gige 4/0/0/0
    !
    !

```

Related Commands

| Command | Description |
|--|--|
| address-family ipv4 unicast (BFD), on page 527 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| bfd fast-detect, on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd multiplier, on page 555 | Sets the BFD multiplier. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

bfd multipath include location

To include specific linecards to host BFD multiple path sessions, use the **bfd multipath include location** command in the global configuration mode. To remove the configuration, use the **no** form of this command.

bfd multipath include location *node-id*

no bfd multipath include location *node-id*

Syntax Description

| | |
|--------------------------------|---|
| location <i>node-id</i> | Configures BFD multipath on the specified location. The <i>node-id</i> variable is mentioned in the <i>rack/slot/module</i> notation. |
|--------------------------------|---|

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|---|
| Release 4.2.0 | This command was introduced. |
| Release 4.3.0 | This command was supported for BFD IPv6 Multihop (BFDv6Mhop) configuration. |
| Release 4.3.1 | Support for this command was removed on ASR 9000 Ethernet Line Card. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

This example shows how to run the **bfd multipath include location** command on a specific location:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# bfd multipath include location 0/5/CPU0
```

Related Commands

| Command | Description |
|---|--|
| bfd, on page 529 | Enters BFD configuration mode. |
| show bfd multipath, on page 589 | Displays information regarding BFD multipath sessions. |

bfd multiplier

To set the Bidirectional Forwarding Detection (BFD) multiplier, use the **bfd multiplier** command in the appropriate configuration mode. To return the router to the default setting, use the **no** form of this command.

bfd multiplier *multiplier*

no bfd multiplier [*multiplier*]

Syntax Description

multiplier

Number of times a packet is missed before BFD declares the neighbor down. The ranges are as follows:

- BGP—2 to 16
- IS-IS—2 to 50
- MPLS-TE—2 to 10
- OSPF and OSPFv3—2 to 50
- PIM—2 to 50

Command Default

The default multiplier is 3.

Command Modes

Router configuration
 Interface configuration
 MPLS-TE configuration
 Router configuration
 Area configuration
 Area interface configuration
 Router configuration
 Area configuration
 Area interface configuration
 Interface configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | The bfd multiplier command was supported in the following modes: OSPFv3 router configuration, OSPFv3 area configuration, and OSPFv3 area interface configuration. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In OSPF and OSPFv3 environments, the setting of the **bfd multiplier** command is inherited from the highest-level configuration mode in which the command was configured. From the lowest to the highest configuration modes, the inheritance rules are as follows:

- If you configure a multiplier in area interface configuration mode, the updated multiplier affects the BFD sessions on the specified interface only.
- If you configure a multiplier in area configuration mode, the updated multiplier affects the BFD sessions on all interfaces in the specified area.
- If you configure a multiplier in router configuration mode, the updated multiplier affects the BFD sessions in all areas and all associated interfaces in the specified routing process.

If desired, you can override these inheritance rules by explicitly configuring the **bfd multiplier** command for a specific area interface or area.

If the multiplier is changed using the **bfd multiplier** command, the new value is used to update all existing BFD sessions for the protocol (BGP, IS-IS, MPLS-TE, OSPF, or OSPFv3).

Task ID

| Task ID | Operations |
|-----------|-------------|
| bgp | read, write |
| isis | read, write |
| mpls-te | read, write |
| ospf | read, write |
| multicast | read, write |

Examples

The following example shows how to set the BFD multiplier in a BGP routing process:

```
RP/0/RSP0/CPU0:router(config)# router bgp 65000
RP/0/RSP0/CPU0:router(config-bgp)# bfd multiplier 2
```

The following example shows the configuration of an OSPFv3 routing process named `san_jose`. The example shows two areas, each of which includes Gigabit Ethernet interfaces. In area 0, the multiplier is set to 5 at the area level, which means that by virtue of the inheritance rules, the same value is set on all interfaces within the area except those on which a different value is explicitly configured. Given this rule, Gigabit Ethernet interface 1/0/0/0 uses a multiplier of 5, which is inherited from the area, while interface 2/0/0/0 uses the explicitly configured multiplier of 2.

In area 1, a multiplier is not configured at the area or interface levels, which means that interfaces 3/0/0/0 and 4/0/0/0 use the default value of 3.

```
router ospfv3 san_jose
bfd fast-detect
  area 0
bfd multiplier 5
int gige 1/0/0/0
  !
int gige 2/0/0/0
bfd multiplier 2
  !
  area 1
int gige 3/0/0/0
  !
int gige 4/0/0/0
  !
  !
```

Related Commands

| Command | Description |
|--|--|
| address-family ipv4 unicast (BFD), on page 527 | Enables BFD fast-detection on a specific IPV4 unicast destination address prefix and on the forwarding next-hop address. |
| bfd fast-detect, on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

clear bfd counters

To clear Bidirectional Forwarding Detection (BFD) counters, use the **clear bfd counters** command in EXEC mode.

clear bfd counters {**ipv4** [**singlehop** | **multihop**] | **ipv6** [**singlehop** | **multihop**] | **all** | **label**} [**packet**] [**timing**] [**interface** *type interface-path-id*] **location** *node-id*

Syntax Description

| | |
|--------------------------------|---|
| ipv4 | (Optional) Clears BFD over IPv4 information only. |
| ipv6 | (Optional) Clears BFD over IPv6 information only. |
| singlehop | (Optional) Clears BFD singlehop information only. |
| multihop | (Optional) Clears BFD multihop information only. |
| all | (Optional) Clears both BFD over IPv4 and BFD over IPv6 information. |
| packet | (Optional) Specifies that packet counters are cleared. |
| timing | (Optional) Specifies that timing counters are cleared. |
| interface | (Optional) Specifies the interface from which the BFD packet counters are cleared. |
| <i>type</i> | Specifies the interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| location <i>node-id</i> | Clears BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | Support for the ipv6 keyword was added. |
| Release 4.2.0 | Support for the singlehop and multihop keywords were added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack/slot/module/port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
 - *rack*: Chassis number of the rack.
 - *slot*: Physical slot number of the line card.
 - *module*: Module number. A physical layer interface module (PLIM) is always 0.
 - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Both IPv4 and IPv6 BFD sessions can run simultaneously on the same line card.

Examples

The following example shows how to clear the BFD IPv6 packet counters on a POS interface:

```
RP/0/RSP0/CPU0:router# clear bfd counters packet ipv6 interface POS 0/1/0/0 location 0/1/cpu0
```

The following example shows how to clear the BFD IPv4 timing counters:

```
RP/0/RSP0/CPU0:router# clear bfd counters ipv4 timing location 0/5/cpu0
```

Related Commands

| Command | Description |
|--|---|
| bfd fast-detect , on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier , on page 555 | Sets the BFD multiplier. |
| show bfd , on page 578 | Displays BFD information for a specific location. |

dampening (BFD)

To specify delays for BFD session startup, use the **dampening** command in Bidirectional Forwarding Detection (BFD) configuration mode. To return to the default, use the **no** form of this command.

dampening [**bundle-member**] {**initial-wait**| **maximum-wait**| **secondary-wait**} *milliseconds*

no dampening [**bundle-member**] {**initial-wait**| **maximum-wait**| **secondary-wait**} *milliseconds*

Syntax Description

| | |
|-----------------------|---|
| bundle-member | (Optional) Specifies initial, maximum, or secondary delays in milliseconds for BFD session startup on BFD bundle members. |
| initial-wait | Specifies the initial delay in milliseconds before starting a BFD session. For bundle members, the default is 16000. For non-bundle interfaces, the default is 2000. |
| maximum-wait | Specifies the maximum delay in milliseconds before starting a BFD session. For bundle members, the default is 600000. For non-bundle interfaces, the default is 12000. Note The maximum delay must be greater than the initial delay. |
| secondary-wait | Specifies a secondary delay in milliseconds before starting a BFD session. For bundle members, the default is 20000. For non-bundle interfaces, the default is 5000. |
| <i>milliseconds</i> | Number from 1 to 3600000. |

Command Default

BFD session startup delays are not configured and the default timer is indefinite.

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.9.0 | This command was introduced. |
| Release 4.0.0 | The bundle-member keyword was added. |
| Release 4.2.0 | The l3-only-mode was added to the other dampening options. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You do not have to configure the BFD startup timers. If you do configure the initial wait startup timer (using the **initial-wait** keyword), then it must be less than the value of the maximum wait timer.

By default, BFD dampening is applied to all sessions in the following manner:

- If a session is brought down, then dampening is applied before a session is allowed to transition to initial/up states.
- Length of time a session is dampened grows exponentially with continuous session flap.
- If a session remains up for minimum two minutes, then the length of time a session dampens with the next session flap is reset to the initial dampening value.

BFD on bundle member applies dampening, only if the detected failure is specific to layer 3. BFD dampening is not invoked for L1 or L2 failures. BFD is started after Layer 1 and Layer 2 (LACP) is up to prevent a race condition and false triggers. BFD is notified to stop/ignore when L1 or L2 goes down and must be notified to start/resume when L1 or L2 recovers for a given/affected link/member.

BFD applies dampening till the session transitions from up to down state and the session is not removed. Whenever there is a failure detected at L1 or L2, the bundle manager removes BFD session on a member.

When dampening is removed a syslog message 'Exponential backoff dampening for BFD session has been cleared for specified BFD session. When/if same session gets created by application(s), only calculated initial wait time will be applied' is generated. If this is the desired behavior, then dampening can be enabled by configuring the BFD configuration, by using the command **bfd dampening bundle-member l3-failure-only**.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to configure an initial and maximum delay for BFD session startup on BFD bundle members:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# dampening bundle-member initial-wait 8000
RP/0/RSP0/CPU0:router(config-bfd)# dampening bundle-member maximum-wait 15000
```

The following example shows how to change the default initial-wait for BFD on a non-bundle interface:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# dampening initial-wait 30000
RP/0/RSP0/CPU0:router(config-bfd)# dampening maximum-wait 35000
```

Related Commands

| Command | Description |
|-----------------------------------|--------------------------------|
| bfd , on page 529 | Enters BFD configuration mode. |

echo disable

To disable echo mode on a router or on an individual interface or bundle, use the **echo disable** command in Bidirectional Forwarding Detection (BFD) configuration mode. To return the router to the default configuration where echo mode is enabled, use the **no** form of this command.

echo disable

no echo disable

Syntax Description This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes BFD configuration
BFD interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you are using BFD with Unicast Reverse Path Forwarding (uRPF), you need to use the **echo disable** command to disable echo mode; otherwise, echo packets are rejected.



Note To enable or disable IPv4 uRPF checking on an IPv4 interface, use the **[no] ipv4 verify unicast source reachable-via** command in interface configuration mode. To enable or disable loose IPv6 uRPF checking on an IPv6 interface, use the **[no] ipv6 verify unicast source reachable-via any** command in interface configuration mode.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |
| | ospf | read, write |
| | isis | read, write |

| Task ID | Operations |
|---------|-------------|
| mpls-te | read, write |

Examples

The following example shows how to disable echo mode on a router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo disable
```

The following example shows how to disable echo mode on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo disable
```

Related Commands

| Command | Description |
|---|---|
| bfd , on page 529 | Enters BFD configuration mode. |
| interface (BFD) , on page 571 | Enters BFD interface configuration mode. |
| ipv4 verify unicast source reachable-via | Enables and disables IPv4 uRPF checking on an IPv4 interface. |
| ipv6 verify unicast source reachable-via any | Enables and disables loose IPv6 uRPF checking on an IPv6 interface. |
| show bfd , on page 578 | Displays BFD information. |

echo ipv4 source

To specify the IP address that you want to use as the source address for BFD echo packets, use the **echo ipv4 source** command in BFD or BFD interface configuration mode. To return to the default, use the **no** form of this command.

echo ipv4 source *ip-address*

no echo ipv4 source *ip-address*

Syntax Description

| | |
|-------------------|---|
| <i>ip-address</i> | 32-bit IPv4 address in dotted-decimal format (A.B.C.D). |
|-------------------|---|

Command Default

The IP address of the output interface, or the IP address in the **router-id** command (if configured), is the default address used for an echo packet when the **echo ipv4 source** command is not configured.

Command Modes

BFD configuration
BFD interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you do not configure the IPv4 source address for echo packets, then BFD uses the IP address of the output interface or the address in the **router-id** command if specified.

You can override the default address for BFD echo packets by specifying an IPv4 source address for echo packets globally for all BFD sessions on the router and at an individual interface. Specifying the IP address at an individual interface will override any value specified globally for BFD on the router.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo ipv4 source 10.10.10.1
```

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets on an individual Gigabit Ethernet interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo ipv4 source 10.10.10.1
```

The following example shows how to specify the IP address 10.10.10.1 as the source address for BFD echo packets on an individual Packet-over-SONET (POS) interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface pos 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# echo ipv4 source 10.10.10.1
```

Related Commands

| Command | Description |
|--|---|
| bfd , on page 529 | Enters BFD configuration mode. |
| echo disable , on page 563 | Disables echo mode on a router or on an individual interface or bundle. |

echo latency detect

To enable latency detection for BFD echo packets, use the **echo latency detect** command in BFD configuration mode. To return to the default, use the **no** form of this command.

echo latency detect [**percentage** *percent-value* [**count** *packet-count*]]

no echo latency detect [**percentage** *percent-value* [**count** *packet-count*]]

Syntax Description

| | |
|--|--|
| percentage <i>percent-value</i> | (Optional) Percentage of the echo failure detection time to be detected as bad latency. The range is 100 to 250. The default is 100. |
| count <i>packet-count</i> | (Optional) Number of consecutive packets received with the detected bad latency that will take down a BFD session. The range is 1 to 10. The default is 1. |

Command Default

Echo latency detection is disabled.

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

Latency detection is only valid where echo mode is supported for BFD. However, it is not supported on bundle interfaces.

Without latency detection, standard BFD echo failure detection tracks only the absence of receipt of echo packets within a period of time based on a counter. However, this standard echo failure detection does not address latency between transmission and receipt of any specific echo packet, which can build beyond desired tolerances over the course of the BFD session.

When latency detection is enabled, a percentage is multiplied to the echo failure detection value ($I \times M \times \%$), and the roundtrip delay is computed for the echo packet. If this delay is greater than ($I \times M \times \%$), then the BFD session is taken down.

If you have specified a packet count, then the system tracks the number of packets received back-to-back with bad latency before taking down the session.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

In the following examples, consider that the BFD minimum interval is 50 ms, and the multiplier is 3 for the BFD session.

The following example shows how to enable echo latency detection using the default values of 100% of the echo failure period (I x M) for a packet count of 1. In this example, when one echo packet is detected with a roundtrip delay greater than 150 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect
```

The following example shows how to enable echo latency detection based on 200% (two times) of the echo failure period for a packet count of 1. In this example, when one packet is detected with a roundtrip delay greater than 300 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect percentage 200
```

The following example shows how to enable echo latency detection based on 100% of the echo failure period for a packet count of 3. In this example, when three consecutive echo packets are detected with a roundtrip delay greater than 150 ms, the session is taken down:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo latency detect percentage 100 count 3
```

Related Commands

| Command | Description |
|---|---|
| bfd , on page 529 | Enters BFD configuration mode. |
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier , on page 555 | Sets the BFD multiplier. |
| echo startup validate , on page 569 | Enables verification of the echo packet path before starting a BFD session. |

echo startup validate

To enable verification of the echo packet path before starting a BFD session, use the **echo startup validate** command in BFD configuration mode. To return to the default, use the **no** form of this command.

echo startup [force]

no echo startup [force]

Syntax Description

force (Optional) Ignores the remote 'Required Min Echo RX Interval' setting.

Command Default

Echo startup validation is disabled.

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

Echo validation is only valid where echo mode is supported for BFD. However, it is not supported on bundle interfaces.

When a BFD session is down and the **echo startup validate** command is configured, an echo packet is periodically transmitted on the link while it is down to verify successful transmission within the configured latency before allowing the BFD session to change state.

Without the **force** option, the echo validation test only runs if the last received control packet contains a non-zero "Required Min Echo RX Interval" value. When the **force** keyword is configured, the echo validation test runs regardless of this value.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |

| Task ID | Operations |
|---------|-------------|
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to enable echo startup validation for BFD sessions on non-bundle interfaces if the last received control packet contains a non-zero “Required Min Echo RX Interval” value:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo startup validate
```

The following example shows how to enable echo startup validation for BFD sessions on non-bundle interfaces regardless of the “Required Min Echo RX Interval” value in the last control packet:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# echo startup validate force
```

Related Commands

| Command | Description |
|---|---|
| bfd , on page 529 | Enters BFD configuration mode. |
| echo latency detect , on page 567 | Enables latency detection for BFD echo packets. |

interface (BFD)

To enter Bidirectional Forwarding Detection (BFD) interface configuration mode, where you can disable echo mode on an interface, use the **interface** command in BFD configuration mode. To return to BFD configuration mode, use the **no** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack/slot/module/port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
 - *rack*: Chassis number of the rack.
 - *slot*: Physical slot number of the line card.
 - *module*: Module number. A physical layer interface module (PLIM) is always 0.
 - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

If you are using BFD with Unicast Reverse Path Forwarding (uRPF) on a particular interface, then you need to use the **echo disable** command in BFD interface configuration mode to disable echo mode on that interface; otherwise, echo packets are rejected by the interface.



Note To enable or disable IPv4 uRPF checking on an IPv4 interface, use the **[no] ipv4 verify unicast source reachable-via** command in interface configuration mode. To enable or disable loose IPv6 uRPF checking on an IPv6 interface, use the **[no] ipv6 verify unicast source reachable-via any** command in interface configuration mode.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to enter BFD interface configuration mode for a Gigabit Ethernet interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)#
```

The following example shows how to enter BFD interface configuration mode for a Packet-over-SONET/SDH (POS) interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface pos 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)#
```

Related Commands

| Command | Description |
|--|--|
| bfd , on page 529 | Enters BFD configuration mode. |
| echo disable , on page 563 | Disables echo mode on an individual interface or on the entire router. |
| ipv4 verify unicast source reachable-via | Enables and disables IPv4 uRPF checking on an IPv4 interface. |
| ipv6 verify unicast source reachable-via any | Enables and disables loose IPv6 uRPF checking on an IPv6 interface. |

| Command | Description |
|--|---------------------------|
| show bfd , on page 578 | Displays BFD information. |

ipv6 checksum

To enable and disable IPv6 checksum calculations on BFD UDP packets globally or on a BFD interface, use the **ipv6 checksum** command in Bidirectional Forwarding Detection (BFD) or BFD interface configuration mode. To return to the default, use the **no** form of this command.

BFD Configuration

ipv6 checksum disable

no ipv6 checksum disable

BFD Interface Configuration

ipv6 checksum [disable]

no ipv6 checksum [disable]

Syntax Description

| | |
|----------------|--|
| disable | (Optional for BFD interface configuration only) Disables IPv6 checksum calculations. |
|----------------|--|

Command Default

IPv6 checksum calculations on BFD UDP packets is disabled.

Command Modes

BFD configuration

BFD interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IPv6 checksum calculations for UDP packets are disabled by default for BFD sessions. You can enable IPv6 checksum support either globally for all BFD sessions, or on an individual interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |

| Task ID | Operations |
|---------|-------------|
| isis | read, write |
| mpls-te | read, write |

Examples

The following example shows how to enable IPv6 checksum calculations for UDP packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# no ipv6 checksum disable
```

The following example shows how to disable IPv6 checksum calculations for UDP packets for all BFD sessions on the router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# ipv6 checksum disable
```

The following example shows how to enable echo mode for BFD sessions on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# ipv6 checksum
```

The following example shows how to disable echo mode for BFD sessions on an individual interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# bfd
RP/0/RSP0/CPU0:router(config-bfd)# interface gigabitethernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-bfd-if)# ipv6 checksum disable
```

Related Commands

| Command | Description |
|---|---|
| bfd , on page 529 | Enters BFD configuration mode. |
| interface (BFD) , on page 571 | Enters BFD interface configuration mode, where you can disable echo mode on an interface. |
| show bfd , on page 578 | Displays BFD information for a specific location. |

multihop ttl-drop-threshold

To specify the maximum time to live (TTL) value for multihop sessions per system, use the **multihop ttl-drop-threshold** command in the BFD configuration mode. To return to the default, use the **no** form of this command.

multihop ttl-drop-threshold *value*

no multihop ttl-drop-threshold *value*

Syntax Description

value Specifies the configurable range of values for TTL. It ranges from 0 to 254.

Command Default

No default behavior or values

Command Modes

BFD configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

This configuration command is only applicable for BFD multihop sessions.

Task ID

| Task ID | Operation |
|---------|-------------|
| bgp | read, write |
| ospf | read, write |
| isis | read, write |
| mpls-te | read, write |

Examples

This example shows how to set the maximum ttl value as 2 using the **multihop-ttl-drop-threshold** command:

```
RP/0/RSP0/CPU0:routerconfigure
RP/0/RSP0/CPU0:routerbfd multihop ttl-drop threshold 2
```

Related Commands

| Command | Description |
|---|--|
| show bfd counters, on page 583 | Displays BFD counter information. |
| show bfd multipath, on page 589 | Displays information regarding BFD multipath sessions. |

show bfd

To display Bidirectional Forwarding Detection (BFD) information for a specific location, use the **show bfd** command in EXEC mode.

```
show bfd [ipv4| [singlehop| multihop ]| ipv6 [singlehop| multihop ]| all[label]interface[destination| source
] [location node-id]
```

Syntax Description

| | |
|--------------------------------|--|
| ipv4 | (Optional) Displays BFD over IPv4 information only. |
| multihop | (Optional) Displays BFD multihop information only. |
| singlehop | (Optional) Displays BFD singlehop information only. |
| ipv6 | (Optional) Displays BFD over IPv6 information only. |
| all | (Optional) Displays both BFD over IPv4 and BFD over IPv6 information. |
| label | (Optional) Displays the BFD label information. |
| interface | Specifies the BFD interface. |
| destination | (Optional) Specifies the destination IPv4 unicast address. |
| source | (Optional) Specifies the source IPv4 unicast address. |
| location <i>node-id</i> | Displays BFD information for the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | Support for the ipv6 keyword was added. |
| Release 4.2.0 | Support for multihop keyword was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| ospf | read |
| isis | read |
| mpls-te | read |

Examples

The following example shows the output from the **show bfd** command:

```
RP/0/RSP0/CPU0:router# show bfd
IPv4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd all** command:

```
RP/0/RSP0/CPU0:router# show bfd all
IPv4:
-----
IPv4 Sessions Up: 20, Down: 0, Unknown/Retry: 2, Total: 22
IPv6:
-----
IPv6 Sessions Up: 128, Down: 2, Unknown/Retry: 1, Total: 131
Label:
-----
Label Sessions Up: 10, Down: 0, Unknown/Retry: 1, Total: 11
```

The following example shows the output from the **show bfd ipv4** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv4
IPv4 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd ipv6** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv6
IPv6 Sessions Up: 0, Down: 0, Total: 0
```

The following example shows the output from the **show bfd ipv4 location** command:

```
RP/0/RSP0/CPU0:router# show bfd ipv6 session detail interface tenGigE 0/0/0/0.100
I/F: TenGigE0/0/0/0.100, Location: 0/0/CPU0
Dest: 1001:1:1:1:1:1:1:2
Src: 1001:1:1:1:1:1:1:1
```

```

State: UP for 0d:1h:37m:3s, number of times UP: 1
Session type: PR/V6/SH
Received parameters:
Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
Required echo rx interval: 0 us, multiplier: 3, diag: None
My discr: 2148076695, your discr: 2148075550, state UP, D/F/P/C/A: 0/0/0/1/0
Transmitted parameters:
Version: 1, desired tx interval: 300 ms, required rx interval: 300 ms
Required echo rx interval: 0 us, multiplier: 3, diag: None
My discr: 2148075550, your discr: 2148076695, state UP, D/F/P/C/A: 0/1/0/1/0
Timer Values:
Local negotiated async tx interval: 300 ms
Remote negotiated async tx interval: 300 ms
Desired echo tx interval: 0 s, local negotiated echo tx interval: 0 us
Echo detection time: 0 us(0 us*3), async detection time: 900 ms(300 ms*3)
Local Stats:
Intervals between async packets:
Tx: Number of intervals=3, min=10 ms, max=2290 ms, avg=824 ms
   Last packet transmitted 5823 s ago
Rx: Number of intervals=6, min=3 ms, max=2 s, avg=461 ms
   Last packet received 5823 s ago
Intervals between echo packets:
Tx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
   Last packet transmitted 0 s ago
Rx: Number of intervals=0, min=0 s, max=0 s, avg=0 s
   Last packet received 0 s ago
Latency of echo packets (time between tx and rx):
Number of packets: 0, min=0 us, max=0 us, avg=0 us
Session owner information:

```

| Client | Desired | | Adjusted | |
|-------------|----------|------------|----------|------------|
| | Interval | Multiplier | Interval | Multiplier |
| ipv6_static | 300 ms | 3 | 300 ms | 3 |
| ipv6_static | 300 ms | 3 | 300 ms | 3 |

```

H/W Offload Info:
H/W Offload capability : Y, Hosted NPU      : 0/0/CPU0/NPU0
Async Offloaded        : Y, Echo Offloaded : N
Async rx/tx            : 7/4

Platform Info:
NPU ID: 0
Async RTC ID          : 1          Echo RTC ID          : 0
Async Feature Mask    : 0x8        Echo Feature Mask    : 0x0
Async Session ID      : 0x3c07     Echo Session ID      : 0x0
Async Tx Key          : 0x3c070801 Echo Tx Key          : 0x0
Async Tx Stats addr   : 0x3f69e800 Echo Tx Stats addr   : 0x0
Async Rx Stats addr   : 0x4069e800 Echo Rx Stats addr   : 0x0

```

Related Commands

| Command | Description |
|--|---|
| bfd fast-detect , on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier , on page 555 | Sets the BFD multiplier. |

show bfd client

To display Bidirectional Forwarding Detection (BFD) client information, use the **show bfd client** command in EXEC mode.

show bfd client [detail]

| Syntax Description | detail (Optional) Specifies detailed client information including number of sessions and client reconnects. | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | Enter the show bfd client command without specifying the detail keyword to display summarized BFD client information. | | | | |
| Command Modes | EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 3.7.2</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 3.7.2 | This command was introduced. |
| Release | Modification | | | | |
| Release 3.7.2 | This command was introduced. | | | | |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | bgp | read |
| | ospf | read |
| | isis | read |
| | mpls-te | read |

Examples

The following example shows the output from the **show bfd client** command:

```
RP/0/RSP0/CPU0:router# show bfd client
Name           Node           Num sessions
-----
```

show bfd client

```

bgp          0/RSP0/CPU0 0
isis        0/RSP0/CPU0 0
isis        0/RSP0/CPU0 0

```

Table 44: show bfd client Field Descriptions

| Field | Description |
|--------------|---|
| Name | Name of the BFD client. |
| Node | Location of the BFD client. |
| Num sessions | Number of active sessions for the BFD client. |

Related Commands

| Command | Description |
|---|---|
| bfd fast-detect, on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 555 | Sets the BFD multiplier. |
| show bfd, on page 578 | Displays BFD information for a specific location. |

show bfd counters

To display Bidirectional Forwarding Detection (BFD) counter information, use the **show bfd counters** command in EXEC mode.

```
show bfd counters [ipv4] [singlehop| multihop] | ipv6 [singlehop| multihop] | all | label] packet [interface
type interface-path-id] location node-id
```

Syntax Description

| | |
|--------------------------------|---|
| ipv4 | (Optional) Displays BFD over IPv4 information only. |
| ipv6 | (Optional) Displays BFD over IPv6 information only. |
| singlehop | (Optional) Displays BFD singlehop information only. |
| multihop | (Optional) Displays BFD multihop information only. |
| all | (Optional) Displays both BFD over IPv4 and BFD over IPv6 information. |
| packet | Specifies that packet counters are displayed. |
| interface | (Optional) Specifies the interface for which to show counters. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| location <i>node-id</i> | Displays BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | Support for the ipv6 keyword was added. |

| Release | Modification |
|---------------|---|
| Release 4.2.0 | Support for the singlehop and multihop keywords were added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack/slot/module/port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
 - *rack*: Chassis number of the rack.
 - *slot*: Physical slot number of the line card.
 - *module*: Module number. A physical layer interface module (PLIM) is always 0.
 - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| ospf | read |
| isis | read |
| mpls-te | read |

Examples

The following example shows the output from the **show bfd counters packet** command for both IPv4 and IPv6:

```
RP/0/RSP0/CPU0:router# show bfd counters packet all interface POS 0/1/0/0 location 0/1/cpu0
Mon Nov  5 08:49:51.950 UTC
IPv4:
-----
POS 0/1/0/0          Recv      Xmit      Echo:      Recv      Xmit
  Async:             520       515           9400      9400

IPv6:
-----
POS 0/1/0/0          Recv      Xmit      Echo:      Recv      Xmit
  Async:             237       237           0         0
```

The following example shows the output from the **show bfd counters packet** command for IPv4:

```
RP/0/RSP0/CPU0:router# show bfd counters ipv4 packet

IPv4 Singlehop:
GigabitEthernet0/0/1/2          Recv      Xmit
  Async:                        4148      4137      Echo: (  47136)    80192
GigabitEthernet0/1/1/2          Recv      Xmit
  Async:                        116876    125756    Echo: ( 2268192)  2301312
Bundle-Ether10                  Recv      Xmit
  Async:                          2          0      Echo:              0          0
Bundle-Ether20                  Recv      Xmit
  Async:                          91         0      Echo:              0          0

IPv4 Multihop: (Src IP/Dst IP/Vrf Id)
33.15.151.4/33.16.151.4/0x12345678  Recv      Xmit
  Async:                          0          570337
```

Table 45: show bfd counters packet Field Descriptions

| Field | Description |
|-------|---|
| Async | Number of asynchronous mode (control) packets that were received or transmitted on the specified interface. |
| Echo | Number of echo packets that were received or transmitted on the specified interface. |

Related Commands

| Command | Description |
|--|---|
| bfd fast-detect , on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval , on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier , on page 555 | Sets the BFD multiplier. |

show bfd mib session

To display IPv4 and IPv6 Bidirectional Forwarding Detection (BFD) MIB session information, use the **show bfd mib session** command in EXEC mode.

show bfd mib session [**location** *node-id*]

Syntax Description

| | |
|--------------------------------|---|
| location <i>node-id</i> | (Optional) Displays all IPv4 and IPv6 BFD MIB session information stored on the specified node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |
|--------------------------------|---|

Command Default

When *node-id* is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.9.0 | This command was introduced. |
| Release 4.0.0 | Support for the display of IPv6 BFD MIB session information was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When *node-id* is not specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the route processor node, is displayed, and this information is populated and updated only after SNMP operations for those BFD MIB sessions are performed.

When *node-id* is specified, information for all IPv4 and IPv6 BFD MIB sessions, stored on the specified node (linecard), is displayed, and this information is updated automatically without SNMP operations being performed.

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| ospf | read |

| Task ID | Operations |
|---------|------------|
| isis | read |
| mpls-te | read |

Examples

The following example displays all IPv4 and IPv6 BFD MIB session information stored on the RP node:

```
RP/0/RSP0/CPU0:router# show bfd mib session
Tue Sep  9 07:49:30.828 PST DST
Local Discr: 327681(0x50001), Remote Discr: 0(0x0)
BFD session: GigabitEthernet0_1_5_2(0x11800c0), 10.27.4.7
  Current State: ADMIN DOWN, Number of Times UP: 0
  Running Version: 0, Last Down Diag: None
  Last Up Time (s.ns): 0.0
  Last Down Time (s.ns): 0.0
  Detection Multiplier: 0
  Desired Min TX Interval: 0
  Required Min RX Interval: 0
  Required Min RX Echo Interval: 0
  Packets in/out: 0/0
  Current Trap Bitmap: 0x0
  Last Time Cached: Not yet cached
```

The following example displays all IPv4 and IPv6 BFD MIB session information stored on 0/1/CPU0:

```
RP/0/RSP0/CPU0:router# show bfd mib session location 0/1/CPU0
Tue Sep  9 07:44:49.190 PST DST
Local Discr: 327681(0x50001), Remote Discr: 0(0x0)
BFD session: GigabitEthernet0_1_5_2(0x11800c0), 10.27.4.7
  Number of times UP: 0
  Last Down Diag: None
  Last Up Time (s.ns): 0.0
  Last Down Time (s.ns): 0.0
  Packets in/out: 0/1140134
```

Table 46: show bfd mib Field Descriptions

| Field | Description |
|--------------------|--|
| date and timestamp | Date and time stamp during which a snapshot of the BFD MIB session information is taken. |
| Local Discr | Local discriminator (in decimal and hexadecimal) that uniquely identifies the BFD MIB session. |
| Remote Discr | Session discriminator (in decimal and hexadecimal) that was chosen by the remote system for the BFD MIB session. |
| BFD session | Index of interface upon which the BFD MIB session is running. Also, neighboring IP address that is monitored with the BFD MIB session. |

| Field | Description |
|-------------------------------|---|
| Current State | Current state of the BFD MIB session. |
| Number of Times UP | Number of times the BFD MIB session has gone into the up state since the router was last rebooted. |
| Running Version | BFD protocol version number in which the BFD MIB session is running. |
| Last Down Diag | Diagnostic value associated with the last time the BFD MIB session went down. |
| Last Up Time (s.ns) | Value of sysUpTime, in <i>seconds.nanoseconds</i> , during which the BFD MIB session last came up. If such an event does not exist, a zero is displayed. |
| Last Down Time (s.ns) | Value of sysUpTime, in <i>seconds.nanoseconds</i> , during which communication was last lost with the neighbor. If such an event does not exist, a zero is displayed. |
| Detection Multiplier | Failure detection multiplier. |
| Desired Min TX Interval | Minimum interval, in microseconds, preferred by the local system when transmitting BFD control packets. |
| Required Min RX Interval | Minimum interval, in microseconds, that the local system supports between received BFD control packets. |
| Required Min RX Echo Interval | Minimum interval, in microseconds, that the local system supports between received BFD echo packets. |
| Packets in/out | Total number of BFD messages received and transmitted for the BFD MIB session. |
| Current Trap Bitmap | Bits that control the trap for the BFD MIB session. A nonzero value indicates that the trap is generated when the next trap event is triggered. |
| Last Time Cached | When information for the BFD MIB session was last cached. Typically, the information is cached when SNMP operations for the BFD MIB session are performed. |

Related Commands

| Command | Description |
|--|-------------|
| show bfd session , on page 593 | |

show bfd multipath

To display information concerning only BFD multipath sessions, use the **show bfd multipath** command in the EXEC mode.

```
show bfd multipath {ipv4| ipv6| label| all} location node-id
```

Syntax Description

| | |
|--------------------------------|--|
| ipv4 | Displays BFD over IPv4 information only. |
| ipv6 | Displays BFD over IPv6 information only. |
| label | Displays BFD label information. |
| all | Displays both BFD over IPv4 and BFD over IPv6 information. |
| location <i>node-id</i> | Displays BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |
| ospf | read |
| isis | read |
| mpls-te | read |

Examples

This example shows the sample output for **show bfd multipath** command:

```
RP/0/RSP0/CPU0:router#show bfd multipath location 0/5/cpu0
```

| Int/Src Addr | Label/Dest Addr | VRF ID | Discr | Node | State |
|--------------|-----------------|------------|-------|----------|-------|
| pw-ether 1 | 10.10.10.10 | 0x00000002 | 0x4 | 0/5/CPU0 | DOWN |
| tunnel-ip 1 | 1.1.1.1 | 0x8 | 0x5 | 0/5/CPU0 | UP |

show bfd neighbor

To display Bidirectional Forwarding Detection (BFD) information to neighbors, use the **show bfd neighbor** command in EXEC mode.

show bfd neighbor {**detail**|**dampened**|**dampening**}

| Syntax Description | | |
|--------------------|-----------------|---|
| | detail | Displays information in detail about the neighbors. |
| | dampened | Displays dampening information about the configured BFD sessions. |
| | dampened | Displays dampening information about the BFD sessions which are currently dampened. |

Command Default None

Command Modes BFD configuration

| Command History | Release | Modification |
|-----------------|-------------|------------------------------|
| | Release 5.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-----------|
| | | |

Examples This example displays the sample output of **show bfd neighbor detail**

```
RP/0/RSP0/CPU0:router#show bfd neighbor detail
```

```
IPv4 Sessions
NeighAddr          LD/RD          RH/RS          State          Int
10.0.0.2           1/0           Down           Down           Et0/0
Session Host: Software
OurAddr: 10.0.0.1
Handle: 1
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 1000000, MinRxInt: 1000000, Multiplier: 3
```

show bfd neighbor

```

Received MinRxInt: 0, Received Multiplier: 0
Holddown (hits): 0(0), Hello (hits): 1000(0)
Rx Count: 0, Rx Interval (ms) min/max/avg: 0/0/0 last: 1257099 ms ago
Tx Count: 0, Tx Interval (ms) min/max/avg: 0/0/0 last: 1257099 ms ago
Elapsed time watermarks: 0 0 (last: 0)
Dampening:  penalty: 0                - not dampened
              flap count: 0          - half-life: 2 sec
              reuse threshold: 2000  - suppress threshold: 3000
              max suppress time: 6 sec
Registered protocols: Stub CEF
Last packet: Version: 1                - Diagnostic: 0
              State bit: AdminDown    - Demand bit: 0
              Poll bit: 0              - Final bit: 0
              C bit: 0
              Multiplier: 0            - Length: 0
              My Discr.: 0              - Your Discr.: 0
              Min tx interval: 0        - Min rx interval: 0
              Min Echo interval: 0

```

This example displays the sample output of **show bfd neighbor dampening**

```
RP/0/RSP0/CPU0:router#show bfd neighbor dampening
```

```

IPv4 Multihop Sessions
NeighAddr[vrf]                LD/RD          RH/RS          State
20.0.0.1                      1/0           Down           Down
Dampening:  penalty: 0                - not dampened
              flap count: 0          - half-life: 2 sec
              reuse threshold: 2000  - suppress threshold: 3000
              max suppress time: 6 sec

```

This example displays the sample output of **show bfd neighbor dampened**

```
RP/0/RSP0/CPU0:router#show bfd neighbor dampened
```

```

IPv4 Sessions
NeighAddr[vrf]                LD/RD          RH/RS          State
20.0.0.1                      1/0           Down           Down
Dampening:  penalty: 4500             - dampened
              flap count: 10          - half-life: 2 sec
              reuse threshold: 2000  - suppress threshold: 3000
              max suppress time: 6 sec

```

Related Commands

| Command | Description |
|---|--|
| bfd dampening , on page 544 | Configures a device to dampen a flapping Bidirectional Forwarding Detection (BFD) session. |

show bfd session

To display Bidirectional Forwarding Detection (BFD) session information, use the **show bfd session** command in EXEC mode.

```
show bfd [ipv4| [singlehop| multihop]] ipv6| [singlehop| multihop]] all| label] session [interface type
interface-path-id [destination ip-address] [detail][in-label]] location node-id
```

Syntax Description

| | |
|-------------------------------|---|
| ipv4 | (Optional) Displays BFD over IPv4 information only. |
| ipv6 | (Optional) Displays BFD over IPv6 information only. |
| singlehop | (Optional) Displays BFD singlehop information only. |
| multihop | (Optional) Displays BFD multihop information only. |
| all | (Optional) Displays both BFD over IPv4 and BFD over IPv6 information. |
| label | (Optional) Displays the MPLS Transport Profile (MPLS-TP) label BFD information only. |
| interface | (Optional) Specifies the interface for which to show information. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| destination ip-address | (Optional) Displays the BFD session destined for the specified IP address. |
| detail | (Optional) Displays detailed session information, including statistics and number of state transitions. |
| in-label | (Optional) Displays the BFD session with a specific incoming MPLS-TP label. |
| location node-id | (Optional) Displays BFD sessions hosted from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

The default is the default address family identifier (AFI) that is set by the **set default-afi** command, IPv4 or IPv6.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.0.0 | Support for the ipv6 keyword was added. |
| Release 4.2.0 | Support for the singlehop and multihop keywords were added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For the *interface-path-id* argument, use the following guidelines:

- If specifying a physical interface, the naming notation is *rack/slot/module/port*. The slash between values is required as part of the notation. An explanation of each component of the naming notation is as follows:
 - *rack*: Chassis number of the rack.
 - *slot*: Physical slot number of the line card.
 - *module*: Module number. A physical layer interface module (PLIM) is always 0.
 - *port*: Physical port number of the interface.
- If specifying a virtual interface, the number range varies, depending on interface type.

**Note**

Only VRF ID is displayed in the summary CLI (such as **show bfd multiple-path**, **show bfd all session**, **show bfd counters**) and VRF name and VRF ID is displayed in the detailed CLI (such as **show bfd all session detail**, **show bfd all session status**).

Task ID

| Task ID | Operations |
|---------|------------|
| bgp | read |
| ospf | read |
| isis | read |
| mpls-te | read |

Examples

The following example shows the output from the **show bfd session** command with the **detail** keyword and IPv4 as the default:

```
RP/0/RSP0/CPU0:router# show bfd session detail

I/f:TenGigE0/2/0/0.6, Location:0/2/CPU0, dest:10.0.6.2, src:10.0.6.1
  State:UP for 0d:0h:3m:4s, number of times UP:1
Received parameters:
  Version:1, desired tx interval:2 s, required rx interval:2 s
  Required echo rx interval:1 ms, multiplier:3, diag:None
  My discr:589830, your discr:590028, state UP, D/F/P/C/A:0/0/0/1/0
Transmitted parameters:
  Version:1, desired tx interval:2 s, required rx interval:2 s
  Required echo rx interval:1 ms, multiplier:3, diag:None
  My discr:590028, your discr:589830, state UP, D/F/P/C/A:0/0/0/1/0
Timer Values:
  Local negotiated async tx interval:2 s
  Remote negotiated async tx interval:2 s
  Desired echo tx interval:250 ms, local negotiated echo tx interval:250 ms
  Echo detection time:750 ms(250 ms*3), async detection time:6 s(2 s*3)
Local Stats:
  Intervals between async packets:
    Tx:Number of intervals=100, min=952 ms, max=2001 ms, avg=1835 ms
      Last packet transmitted 606 ms ago
    Rx:Number of intervals=100, min=1665 ms, max=2001 ms, avg=1828 ms
      Last packet received 1302 ms ago
  Intervals between echo packets:
    Tx:Number of intervals=100, min=250 ms, max=252 ms, avg=250 ms
      Last packet transmitted 188 ms ago
    Rx:Number of intervals=100, min=250 ms, max=252 ms, avg=250 ms
      Last packet received 187 ms ago
  Latency of echo packets (time between tx and rx):
    Number of packets:100, min=1 ms, max=2 ms, avg=1 ms
Session owner information:
  Client          Desired interval      Multiplier
  -----
  bgp-            250 ms                  3
```

The following example shows the output from the **show bfd session** command with the **all** keyword, which displays both IPv4 and IPv6 information:

```
RP/0/RSP0/CPU0:router# show bfd all session location 0/1/CPU0

Mon Nov  5 08:51:50.339 UTC
IPv4:
-----
Interface          Dest Addr          Local det time(int*mult)  State
                   Echo              Async
-----
PO0/1/0/0          10.0.0.2           300ms(100ms*3)          6s(2s*3)          UP

IPv6:
-----
Interface          Dest Addr          Local det time(int*mult)  State
                   Echo              Async
-----
PO0/1/0/0          abcd::2           0s(0s*0)                15s(5s*3)          UP
```

Table 47: show bfd session detail command Field Descriptions

| Field | Description |
|---------------------|---|
| I/f | Interface type. |
| Location | Location of the node that hosts the local endpoint of the connection, in the <i>rack/slot/module</i> notation |
| dest | IP address of the destination endpoint. |
| src | IP address of the source endpoint. |
| State | Current state of the connection, and the number of days, hours, minutes, and seconds that this connection has been active. |
| number of times UP | Number of times this connection has been brought up. |
| Received parameters | <p>Provides information on the last transmitted control packet for the session:</p> <ul style="list-style-type: none"> • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval. • Required echo rx interval—Required echo receive interval. • multiplier— Number of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the peer system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown. |

| Field | Description |
|------------------------|--|
| Transmitted parameters | <p>Provides information on the last transmitted control packet for the session:</p> <ul style="list-style-type: none"> • Version—Version number of the BFD protocol. • desired tx interval—Desired transmit interval. • required rx interval—Required receive interval • Required echo rx interval—Required echo receive interval • multiplierNumber of times a packets is missed before BFD declares the neighbor down. • diag—diagnostic code specifying the local system's reason for the last transition of the session from Up to some other state. • My discr—unique, nonzero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. • your discr— discriminator received from the corresponding remote system. This field reflects back the received value of My discr, or is zero if that value is unknown. |

| Field | Description |
|--------------|--|
| Timer Values | <p data-bbox="922 285 1482 348">Provides information on the timer values used by the local and remote ends, as follows:</p> <ul data-bbox="966 369 1482 1094" style="list-style-type: none"><li data-bbox="966 369 1482 464">• Local negotiated async tx interval—interval at which control packets are being transmitted by the local end.<li data-bbox="966 485 1482 579">• Remote negotiated async tx interval—interval at which control packets should be transmitted by the remote end.<li data-bbox="966 600 1482 663">• Desired echo tx interval—interval at which the local end would like to transmit echo packets.<li data-bbox="966 684 1482 779">• local negotiated echo tx interval—interval at which echo packets are being transmitted by the local end.<li data-bbox="966 800 1482 915">• Echo detection time—local failure detection time of echo packets. It is the product of the local negotiated echo tx interval and the local multiplier.<li data-bbox="966 936 1482 1094">• async detection time—local failure detection time of the asynchronous mode (control packets). It is the product of the remote negotiated async tx interval and the remote multiplier. |

| Field | Description |
|-------------|---|
| Local Stats | <p>Displays the local transmit and receive statistics,</p> <ul style="list-style-type: none"> • Intervals between async packets—provides measurements on intervals between control packets (tx and rx): <ul style="list-style-type: none"> ◦ Number of intervals—number of sampled intervals between control packets ◦ min—minimum measured interval between 2 consecutive control packets ◦ max—maximum measured interval between 2 consecutive control packets ◦ avg—average measured interval between 2 consecutive control packets ◦ Last packet received/transmitted—indicates how long ago the last control packet was received/transmitted. • Intervals between echo packets—provides measurements on intervals between echo packets (tx and rx). The measurements have the same meaning as for async packets. • Latency of echo packets (time between tx and rx)—provides measurements on latency of echo packets, i.e. the time between tx and rx of echo packets: <ul style="list-style-type: none"> ◦ Number of packets—number of sampled echo packets. ◦ min—minimum measured latency for echo packets. ◦ max—maximum measured latency of echo packets. ◦ avg—average measured latency of echo packets. |

| Field | Description |
|---------------------------|--|
| Session owner information | <p>Provides the following information about the session owner.</p> <ul style="list-style-type: none"> • Client—name of the client application process. • Desired interval—desired interval provided by the client, in milliseconds. • Multiplier—multiplier value provided by the client. |

Related Commands

| Command | Description |
|---|---|
| bfd fast-detect, on page 546 | Enables BFD to detect failures in the path between adjacent forwarding engines. |
| bfd minimum-interval, on page 549 | Specifies the minimum control packet interval for BFD sessions for the corresponding BFD configuration scope. |
| bfd multiplier, on page 555 | Sets the BFD multiplier. |
| show bfd mib session, on page 586 | Displays BFD MIB session information. |

show bfd summary

To display the percentage of PPS rate in use per line card, maximum usage of PPS, and total number of sessions, use the **show bfd summary** command in the EXEC mode.

show bfd summary [**private**]**location***node-id*

Syntax Description

| | |
|--------------------------------|--|
| private | Displays the private information. |
| location <i>node-id</i> | Displays BFD counters from the specified location. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| bgp | read |
| ospf | read |
| isis | read |
| mpls-te | read |

Examples

This example shows the sample output from the **show bfd summary** command for a specified location:

```
RP/0/RSP0/CPU0:router#show bfd summary location 0/1/cpu0
```

| Node | PPS % | rate Used | usage Max | Session Total | number Max |
|----------|-------|-----------|-----------|---------------|------------|
| 0/1/CPU0 | 0 | 80 | 9600 | 4 | 4000 |

This example shows the sample output from the **show bfd summary** command:

```
RP/0/RSP0/CPU0:router#show bfd summary
```

| Node | PPS % | rate Used | usage Max | Session Total | number Max |
|----------|-------|-----------|-----------|---------------|------------|
| 0/0/CPU0 | 0 | 0 | 9600 | 0 | 4000 |
| 0/1/CPU0 | 0 | 0 | 9600 | 0 | 4000 |
| 0/2/CPU0 | 0 | 0 | 9600 | 0 | 4000 |
| 0/5/CPU0 | 0 | 0 | 9600 | 0 | 4000 |
| 0/6/CPU0 | 0 | 0 | 9600 | 0 | 4000 |
| 0/7/CPU0 | 0 | 0 | 9600 | 0 | 4000 |



Enhanced Interior Gateway Routing Protocol Commands

This module describes the commands used to configure and monitor the Enhanced Interior Gateway Routing Protocol (EIGRP).

For detailed information about EIGRP concepts, configuration tasks, and examples, see *Implementing EIGRP on Cisco ASR 9000 Series Router* in *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family \(EIGRP\)](#), page 605
- [authentication keychain](#), page 607
- [auto-summary \(EIGRP\)](#), page 609
- [autonomous-system](#), page 611
- [bandwidth-percent \(EIGRP\)](#), page 613
- [clear eigrp neighbors](#), page 615
- [clear eigrp topology](#), page 617
- [default-information](#), page 619
- [default-metric \(EIGRP\)](#), page 621
- [distance \(EIGRP\)](#), page 623
- [hello-interval \(EIGRP\)](#), page 625
- [hold-time \(EIGRP\)](#), page 626
- [interface \(EIGRP\)](#), page 628
- [log-neighbor-changes](#), page 630
- [log-neighbor-warnings](#), page 631
- [maximum-paths \(EIGRP\)](#), page 632
- [maximum-prefix \(EIGRP\)](#), page 634
- [metric \(EIGRP\)](#), page 637

- [metric maximum-hops, page 639](#)
- [metric rib-scale, page 641](#)
- [metric weights, page 643](#)
- [neighbor \(EIGRP\), page 646](#)
- [neighbor maximum-prefix, page 648](#)
- [next-hop-self disable, page 651](#)
- [nsf disable \(EIGRP\), page 652](#)
- [passive-interface \(EIGRP\), page 654](#)
- [redistribute \(EIGRP\), page 655](#)
- [redistribute maximum-prefix, page 658](#)
- [route-policy \(EIGRP\), page 660](#)
- [router eigrp, page 662](#)
- [router-id \(EIGRP\), page 664](#)
- [show eigrp accounting, page 665](#)
- [show eigrp interfaces, page 668](#)
- [show eigrp neighbors, page 673](#)
- [show eigrp topology, page 677](#)
- [show eigrp traffic, page 681](#)
- [show protocols \(EIGRP\), page 683](#)
- [site-of-origin \(EIGRP\), page 687](#)
- [split-horizon disable \(EIGRP\), page 689](#)
- [stub \(EIGRP\), page 690](#)
- [summary-address \(EIGRP\), page 693](#)
- [timers active-time, page 695](#)
- [timers nsf route-hold \(EIGRP\), page 697](#)
- [variance, page 699](#)
- [vrf \(EIGRP\), page 701](#)

address-family (EIGRP)

To enable an IPv4 address family under Enhanced Interior Gateway Routing Protocol (EIGRP), use the **address-family** command in the appropriate mode. To remove the address family from the EIGRP configuration, use the **no** form of this command.

address-family {ipv4| ipv6}

no address-family {ipv4| ipv6}

Syntax Description

| | |
|------|------------------------------|
| ipv4 | Selects IPv4 address family. |
|------|------------------------------|

Command Default

No default behavior or values

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family ipv4** command to configure IPv4 address family sessions in EIGRP.

EIGRP Virtual Private Networks (VPNs) can be configured under IPv4 address family configuration modes after first entering VRF configuration mode. All commands in address family configuration mode can be configured in VRF address families except the **autonomous-system** and **maximum-prefix** commands.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to configure an IPv4 VRF address family session after defining the VRF named vrf1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
```

```
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-metric 1000 100 255 1 1500
```

Related Commands

| Command | Description |
|-------------------------------|---|
| autonomous-system | Configures an EIGRP routing process to run within a VRF. |
| maximum-prefix (EIGRP) | Limits the number of prefixes that are accepted under an address family by EIGRP. |

authentication keychain

To authenticate all EIGRP protocol traffic on one or more interfaces based on the MD5 algorithm, use the **authentication keychain** command in an appropriate configuration mode. To disable authentication, use the **no** form of this command.

authentication keychain *key-chain-name*

no authentication keychain *key-chain-name*

Syntax Description

| | |
|-----------------------|-------------------------------------|
| <i>key-chain-name</i> | Name of the authentication keychain |
|-----------------------|-------------------------------------|

Command Default

Authentication is disabled.

Command Modes

IPv4 address family interface configuration
 IPv6 address family interface configuration
 IPv4 VRF address family interface configuration
 IPv6 VRF address family interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family ipv4** command to configure IPv4 address family sessions in EIGRP and the **address-family ipv6** command to configure IPv6 address family sessions in EIGRP.

EIGRP Virtual Private Networks (VPNs) can be configured under IPv4 and IPv6 address family configuration modes after first entering VRF configuration mode.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to enable an EIGRP authentication keychain:

```
RP/0/RSP0/CPU0:router# configure eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface POS 0/2/0/0  
RP/0/RSP0/CPU0:router(config-eigrp-if)# authentication key chain key1
```

Related Commands

| Command | Description |
|---|--|
| router eigrp, on page 662 | Configures a routing process and enter router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP). |

auto-summary (EIGRP)

To allow automatic summarization of subnet routes into network-level routes for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **auto-summary** command in the appropriate configuration mode. To disable this function and send subprefix routing information across classful network boundaries, use the **no** form of this command.

auto-summary

no auto-summary

Syntax Description This command has no keywords or arguments.

Command Default The behavior of this command is disabled by default. (The software sends subnet routing information across classful network boundaries.)

Command Modes IPv4 Address family configuration
IPv4 VRF address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Route summarization reduces the amount of routing information in the routing tables. Use the **auto-summary** command to allow the software to create summary subprefixes to the classful network boundary when crossing classful network boundaries.

EIGRP summary routes are given an administrative distance value of 5. You cannot configure this value.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples The following example shows how to enable automatic summarization for EIGRP 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
```

```
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# auto-summary
```

Related Commands

| Command | Description |
|--|---|
| summary-address (EIGRP), on page 693 | Configures a summary aggregate address for a specified interface. |

autonomous-system

To configure the autonomous-system number for an address-family of Enhanced Interior Gateway Routing Protocol (EIGRP) routing process, use the **autonomous-system** command in the address family configuration mode. To remove the autonomous-system number for an address-family of EIGRP routing process, use the **no** form of this command.

autonomous-system *as-number*

no autonomous-system *as-number*

Syntax Description

| | |
|------------------|--|
| <i>as-number</i> | Autonomous system number of the EIGRP routing process. Range is from 1 to 65535. |
|------------------|--|

Command Default

None

Command Modes

IPv4 address family configuration (For EIGRP configured using virtual-name only).
IPv4 VRF address family configuration.

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 5.1 | This command was added in address family configuration mode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **autonomous-system** command in IPv4 VRF address family configuration mode.

The same VRF-autonomous-system combination cannot be used across multiple process instances.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

This example shows how to configure autonomous system 101 under VRF VPN-1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# autonomous-system 101
```

This example shows how to configure autonomous-system 1 for default/global VRF of EIGRP configured using virtual-name:

```
RP/0/RSP0/CPU0:router(config)# router eigrp name
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# autonomous-system 1
```

Related Commands

| Command | Description |
|--|--|
| vrf (EIGRP), on page 701 | Defines a VRF instance and to enters VRF configuration mode. |

bandwidth-percent (EIGRP)

To configure the percentage of bandwidth that may be used by EIGRP on an interface, use the **bandwidth-percent** command in interface configuration mode. To restore the default value, use the **no** form of this command.

bandwidth-percent *percent*

no bandwidth-percent

| Syntax Description | percent | Percentage of bandwidth that EIGRP may use. |
|--------------------|---------|---|
|--------------------|---------|---|

Command Default *percent* : 50

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

EIGRP uses up to 50 percent of the bandwidth of a link, as defined by the bandwidth interface configuration command. This command may be used if some other fraction of the bandwidth is desired. Values greater than 100 percent may be configured. The configuration option may be useful if the bandwidth is set artificially low for other reasons.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples The following example shows how to configure EIGRP to use up to 75 percent (42 kbps) of an interface in autonomous system 209:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
```

bandwidth-percent (EIGRP)

```
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# bandwidth-percent 75
```

Related Commands

| Command | Description |
|-----------------------|--|
| bandwidth (interface) | Sets a bandwidth value for an interface. |

clear eigrp neighbors

To remove and re-establish Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor entries from the appropriate table, use the **clear eigrp neighbors** command in EXEC configuration mode.

```
clear eigrp [as-number] [vrf {vrf | all}] [ipv4|ipv6] neighbors [ip-address] type interface-path-id [soft]
```

Syntax Description

| | |
|--|---|
| <i>as-number</i> | (Optional) Autonomous system number. Range is from 1 to 65535. |
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| ipv4 | (Optional) Specifies the IPv4 address family. |
| <i>ip-address</i> | (Optional) Address of the neighbor. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| soft | (Optional) Specifies a soft reset. |

Command Default

When no autonomous system number or VRF instance is specified, all EIGRP neighbor entries are cleared from the table.

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to clear all EIGRP VRF entries for neighbor Gigabit Ethernet interface 0/5/0/0:

```
RP/0/RSP0/CPU0:router# clear eigrp vrf customer_1 neighbors GigabitEthernet 0/5/0/0
```

Related Commands

| Command | Description |
|--|---|
| show eigrp interfaces, on page 668 | Displays information about interfaces configured for EIGRP. |
| show eigrp neighbors, on page 673 | Displays the neighbors discovered by EIGRP. |

clear eigrp topology

To remove and relearn Enhanced Interior Gateway Routing Protocol (EIGRP) topology entries from the appropriate table, use the **clear eigrp topology** command in EXEC configuration mode.

```
clear eigrp [as-number] [vrf{vrf | all}] [ipv4 | ipv6] topology [prefix mask | prefix/length]
```

Syntax Description

| | |
|--|--|
| as-number | (Optional) Autonomous system number. Range is from 1 to 65535. |
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| ipv4 | (Optional) Specifies the IPv4 address family. |
| prefix | IP prefix, which limits output to a specific route. |
| mask | IP address mask. |
| / length | Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If <i>length</i> is used, the slash is required. |

Command Default

No EIGRP topology entries are cleared.

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to clear EIGRP topology entries for a specific route:

```
RP/0/RSP0/CPU0:router# clear eigrp topology 10.1.0.0/8
```

Related Commands

| Command | Description |
|--|--|
| show eigrp topology, on page 677 | Displays information for the EIGRP topology table. |

default-information

To control the candidate default routing information for an Enhanced Interior Gateway Routing Protocol (EIGRP), use the **default-information** command in the appropriate configuration mode. To suppress EIGRP candidate default information in incoming or outgoing updates, use the **no** form of this command.

default-information allowed {in| out} [*route-policy name*]

no default-information allowed {in| out} [*route-policy name*]

Syntax Description

| | |
|---------------------------------|--|
| allowed | Specifies EIGRP to allow default routing information. |
| in | Specifies EIGRP to allow inbound default routing information. |
| out | Specifies EIGRP to allow outbound default routing information. |
| route-policy <i>name</i> | (Optional) Specifies a route policy. |

Command Default

Default routing information is not accepted or flagged.

Command Modes

Address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to configure inbound default routes specified with route policy acme to be accepted by an EIGRP peer in autonomous system 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-information accept in route-policy
acme
```

default-metric (EIGRP)

To set metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP), use the **defaultmetric** command in the appropriate configuration mode. To remove the metric values and restore the default state, use the **no** form of this command.

default-metric *bandwidth delay reliability loading mtu*

no default-metric

Syntax Description

| | |
|--------------------|--|
| <i>bandwidth</i> | Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295. |
| <i>delay</i> | Route delay in ten microsecond units. Range is 1 to 4294967295. |
| <i>reliability</i> | Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100-percent reliability; 0 means the link is not reliable. |
| <i>loading</i> | Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100-percent loading). |
| <i>mtu</i> | Minimum maximum transmission unit (MTU) size of the route in bytes. Range is from 1 to 65535. |

Command Default

No default values

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command to provide default metric values while redistributing a protocol into EIGRP.

Metric defaults have been carefully set to work for a wide variety of networks. Take great care when changing these values.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to take redistributed Routing Information Protocol (RIP) metrics and translate them into EIGRP metrics with values as follows: bandwidth = 1000, delay = 100, reliability = 250, loading = 100, and MTU = 1500.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrfl
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute rip
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# default-metric 1000 100 250 100 1500
```

Related Commands

| Command | Description |
|--------------|---|
| redistribute | Redistributes routes from one routing domain into another routing domain. |

distance (EIGRP)

To allow the use of one of two administrative distances—internal and external—that could provide a better route to a node, use the **distance** command in the appropriate configuration mode. To reset these values to their defaults, use the **no** form of this command.

distance *internal-distance external-distance*

no distance

Syntax Description

| | |
|--------------------------|--|
| <i>internal-distance</i> | Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255. |
|--------------------------|--|

| | |
|--------------------------|--|
| <i>external-distance</i> | Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a source external to the AS. The distance can be a value from 1 to 255. |
|--------------------------|--|

Command Default

internal-distance : 90

external-distance : 170

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An administrative distance is a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers. Numerically, an administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored.

Use the **distance** command if another protocol is known to provide a better route to a node than was actually learned through the external Enhanced Interior Gateway Routing Protocol (EIGRP) or some internal routes should be preferred by EIGRP.

To display the default administrative distance for a specified routing process, use the **show protocols EXEC** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set the administrative distance of all EIGRP 1 internal routes (within vrf vpn-1) to 80 and all EIGRP external routes to 130:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# distance 80 130
```

Related Commands

| Command | Description |
|---|--|
| show protocols (EIGRP), on page 683 | Displays information about the Enhanced Interior Gateway Routing Protocol (EIGRP) running on the router. |

hello-interval (EIGRP)

To configure the hello interval for an interface, use the **hello-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-interval *seconds*

no hello-interval

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Hello interval (in seconds). Range is from 1 to 65535. |
|----------------|--|

Command Default

For low-speed, nonbroadcast multiaccess (NBMA) networks: 60 seconds

For all other networks: 5 seconds

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set the hello interval to 10 seconds for the interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hello-interval 10
```

hold-time (EIGRP)

To configure the hold time for an interface, use the **hold-time** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hold-time *seconds*

no hold-time

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Hold time (in seconds). Range is from 1 to 65535. |
|----------------|---|

Command Default

Three times the default hello interval time of 15 seconds.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On very congested and large networks, the default hold time might not be sufficient time for all routers to receive hello packets from their neighbors. In this case, you may want to increase the hold time.

We recommend that the hold time be at least three times the hello interval. If a router does not receive a hello packet within the specified hold time, routes through this router are considered unavailable.

Increasing the hold time delays route convergence across the network.

To ensure nonstop forwarding during RP failovers, we recommend that the hold time be increased to 30 seconds.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set the hold time to 0 to 40 seconds for the interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
```

```
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1  
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hold-time 40
```

Related Commands

| Command | Description |
|---|--|
| bandwidth (interface) | Sets a bandwidth value for an interface. |
| hello-interval (EIGRP), on page 625 | Configures the hello interval for the EIGRP routing process designated by an autonomous system number. |

interface (EIGRP)

To define the interfaces on which the Enhanced Interior Gateway Routing Protocol (EIGRP) routing protocol runs, use the **interface** command in the appropriate configuration mode. To disable EIGRP routing for interfaces, use the **no** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

When you do not specify this command in configuration mode, EIGRP routing for interfaces is not enabled.

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an EIGRP process. The interface remains associated with the process even when the IPv4 address of the interface changes.

This command places the router in interface configuration mode, from which you can configure interface-specific settings. Commands configured under this mode (such as the **hello-interval** command) are automatically bound to that interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to enter interface configuration mode for EIGRP process 1 and set the hello interval to 10 seconds for GigabitEthernet interface 0/1/0/0:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp)# router-id 10.1.1.1  
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# hello-interval 10
```

log-neighbor-changes

To enable the logging of changes in Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor adjacencies, use the **log-neighbor-changes** command in the appropriate configuration mode. To disable the logging of changes in EIGRP neighbor adjacencies, use the **no** form of this command.

log-neighbor-changes

no log-neighbor-changes

Command Default

This command has no keywords or arguments.

Adjacency changes are not logged.

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log-neighbor-changes** command to log neighbor adjacency changes, monitor the stability of the routing system, and help detect problems. Logging is disabled by default. To disable the logging of neighbor adjacency changes, use the **no** form of this command.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to enable logging of neighbor changes for EIGRP 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# log-neighbor-changes
```

log-neighbor-warnings

To enable the logging of Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor warning messages, use the **log-neighbor-warnings** command in the appropriate configuration mode. To disable the logging of EIGRP neighbor warning messages, use the **no** form of this command.

log-neighbor-warnings

no log-neighbor-warnings

Syntax Description This command has no keywords or arguments.

Command Default Neighbor warning messages are not logged.

Command Modes IPv4 address family configuration
IPv4 VRF address family configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log-neighbor-warnings** command to disable and enable neighbor warning messages. When neighbor warning messages occur, they are not logged by default.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | eigrp | read, write |

Examples The following example shows how to configure log neighbor warning messages for EIGRP process 20:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 20
RP/0/RSP0/CPU0:router(config-eigrp) vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf) # address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af) # log-neighbor-warnings
```

maximum-paths (EIGRP)

To control the maximum number of parallel routes that the Enhanced Interior Gateway Routing Protocol (EIGRP) can support, use the **maximum-paths** command in the appropriate configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum-paths *maximum*

no maximum-paths

Syntax Description

| | |
|----------------|---|
| <i>maximum</i> | Maximum number of parallel routes that EIGRP can install in a routing table. Range is from 1 to 32 routes . |
|----------------|---|

Command Default

maximum : 4

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-paths** command to allow the EIGRP protocol to install multiple paths into the routing table for each prefix. Multiple paths are installed for both internal and external routes, providing these routes are learned in the same autonomous system and that they are equal cost (according to the EIGRP best path algorithm).

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to allow a maximum of 10 paths to a destination:

```
RP/RSP0/CPU0:router(config)# router eigrp 1
```

```
RP/0/RSP0/CPU0:router(config-eigrp) vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# maximum-paths 10
```

maximum-prefix (EIGRP)

To limit the number of prefixes that are accepted under a VRF address family by Enhanced Interior Gateway Routing Protocol (EIGRP), use the **maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

maximum-prefix *maximum* [*threshold*] [**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] [**warning-only**]

no maximum-prefix *maximum* [*threshold*] [**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*] [**warning-only**]

Syntax Description

| | |
|------------------------------------|--|
| <i>maximum</i> | Maximum number of prefixes allowed under an address family. Range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router. |
| <i>threshold</i> | (Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent. |
| dampened | (Optional) A decay penalty is applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes. This keyword is disabled by default. |
| reset-time <i>minutes</i> | (Optional) The restart count is reset to 0 after the default or user-defined reset time period has expired. The range of values that can be applied with the <i>minutes</i> argument is from 1 to 65535 minutes. The default reset-time period is 15 minutes. |
| restart <i>minutes</i> | (Optional) A time period when router adjacencies are not formed or when redistributed routes are not accepted from the RIB after the maximum prefix limit has been exceeded. The value for the <i>minutes</i> argument is from 1 to 65535 minutes. The default restart time period is 5 minutes. |
| restart-count <i>number</i> | (Optional) Number of times a peering session is automatically reestablished after the peering session is torn down or after the redistribute route is cleared and relearned when the maximum prefix exceeds limits. The default restart count limit is 3. Once the restart count threshold is crossed, you need to use the clear route or clear eigrp neighbors command to reestablish normal peering and redistribution. |
| warning-only | (Optional) Configures the router to generate syslog messages only when the maximum prefix limit is reached, instead of terminating the peering session. |

Command Default

threshold: 75 percent

dampened : False
reset-time : 15 minutes
restart : 5 minutes
restart-count : 3
warning-only : False

Command Modes IPv4 VRF address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-prefix** command to limit the number of prefixes that are accepted from all sources. When the maximum prefix limit is exceeded, sessions with remote peers are torn down, all routes learned from remote peers and through redistribution are removed from the topology and routing tables, and redistribution and peering are suspended for the default or user-defined time period.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples The following example shows how to configure the maximum prefix limit for an EIGRP process, which includes routes learned through redistribution and routes learned through EIGRP peering sessions. The maximum limit is set to 50,000 prefixes. When the number of prefixes learned through redistribution reaches 37,500 (75 percent of 50,000), warning messages are displayed in the console. When the maximum prefix limit is exceeded, all peering sessions are reset, the topology and routing tables are cleared and redistributed routes and all peering sessions are placed in a penalty state.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# maximum-prefix 50000
```

| Related Commands | Command | Description |
|------------------|---|--|
| | clear eigrp neighbors , on page 615 | Deletes EIGRP VPN neighbor entries from the table. |

| Command | Description |
|-------------|---|
| clear route | Deletes routes from the IP routing table. |

metric (EIGRP)

To set metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **metric** command in interface configuration mode. To remove the metric values and restore the default state, use the **no** form of this command.

metric {**bandwidth**| **delay** [**picoseconds**]| **load**| **reliability**}

no metric

Syntax Description

| | |
|--------------------|--|
| bandwidth | Minimum interface bandwidth of the route in kilobits per second. Range is 1 to 4294967295. |
| delay | Interface route delay in tens of microseconds. Delay is 1 or any positive number that is a multiple of 39.1 nanoseconds. Range is 1 to 4294967295. Use the picoseconds keyword to specify interface route delay in picoseconds. <ul style="list-style-type: none"> If picoseconds is not used, the interface route delay in tens of microsecond (default). Range is 1 to 4294967295 If picoseconds is used, interface route delay in picoseconds. Range is 1 to 4294967295 |
| picoseconds | (Optional) Specifies interface route delay in picosends. |
| load | Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100-percent loaded). |
| reliability | Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100-percent reliability; 0 means no reliability. |

Command Default

Metric values are not set.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.6.2 | This command was introduced. |
| Release 4.3.0 | The picoseconds keyword was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **metric** command to provide metric values while redistributing a protocol into an EIGRP interface. Metric defaults have been carefully set to work for a wide variety of networks. Take great care when changing these values.

The picoseconds keyword is supported only in 64 bit mode.

Task ID

| Task ID | Operation |
|---------|-------------|
| eigrp | read, write |

Examples

This example shows how to configure metric values for interface POS 0/1/0/1 with values as bandwidth = 100, delay = 7, reliability = 250, and load = 100.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface POS 0/1/0/1
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# metric bandwidth 100 delay 7 reliability 250
load 100
```

This example shows how to set the delay of interface GigabitEthernet0/0/0/0 to 100 picoseconds:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)#address-family ipv6
RP/0/RSP0/CPU0:router(config-eigrp-af-if)#metric delay 100 picoseconds
```

Related Commands

| Command | Description |
|---|---|
| default-metric (EIGRP), on page 621 | Sets metrics for an Enhanced Interior Gateway Routing Protocol (EIGRP). |

metric maximum-hops

To advertise as unreachable those Enhanced Interior Gateway Routing Protocol (EIGRP) routes with a hop count higher than is specified by the command, use the **metric maximum-hops** command in the appropriate configuration mode. To reset the value to the default, use the **no** form of this command.

metric maximum-hops *hops-number*

no metric maximum-hops

| | | |
|---------------------------|--------------------|---|
| Syntax Description | hops-number | Maximum hop count. Range is from 1 to 255 hops. |
|---------------------------|--------------------|---|

Command Default *hops-number* : 100

Command Modes
IPv4 address family configuration
IPv4 VRF address family configuration

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines
To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **metric maximum-hops** command to provide a safety mechanism that breaks any potential *count-to-infinity* problems. This command causes the EIGRP routing protocol to advertise as unreachable routes with a hop count greater than the value assigned to the *hops-number* argument.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | eigrp | read, write |

Examples
The following example shows how to configure a hop count to 200 for a router that has a complex WAN generating a large hop count under normal (nonlooping) operations.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# metric maximum-hops 200
```

Related Commands

| Command | Description |
|--|---|
| metric weights , on page 643 | Allows the tuning of the EIGRP metric calculations. |

metric rib-scale

To set a RIB scale for EIGRP, use the **metric rib-scale** command in an appropriate configuration mode. To remove the RIB scale and restore the default state, use the **no** form of this command.

metric rib-scale *scale-down-value*

no metric rib-scale

Syntax Description

| | |
|-------------------------|--|
| <i>scale-down-value</i> | Amount to divide the EIGRP Wide Metric by to convert to a 4 Byte RIB metric. Legal ranges are 1-256. Results are expressed as whole integers only (no rounding). |
|-------------------------|--|

Command Default

Default RIB scale is 128 in the 64 bit mode.

In 32 bit mode, rib scale is always 1.

Command Modes

IPv4 address family configuration

IPv6 address family configuration

IPv4 VRF address family configuration

IPv6 VRF address family configuration (Only supported in 64 bit mode)

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The ability to scale the RIB metric is provided through the use of the **metric rib-scale** configuration command. When entered, the command will result in all routes in the RIB being cleared and replaced with the new metric values.

Task ID

| Task ID | Operation |
|---------|-------------|
| eigrp | read, write |

Examples

This example shows how to set the metric rib-scale as 64:

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#router eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp)#address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af)#metric rib-scale 64
```

metric weights

To allow the tuning of the Enhanced Interior Gateway Routing Protocol (EIGRP) metric calculations, use the **metric weights** command in the appropriate configuration mode. To reset the values to their defaults, use the **no** form of this command.

metric weights *tos k1 k2 k3 k4 k5 k6*

no metric weights

Syntax Description

| | |
|--------------------------|---|
| <i>tos</i> | Type of service (ToS) which must always be 0. |
| <i>k1 k2 k3 k4 k5 k6</i> | Constants that convert an EIGRP metric vector into a scalar quantity. The range is 0 to 4294967295. |

Command Default

tos: 0
k1: 1
k2: 0
k3: 1
k4: 0
k5: 0
k6: 0

Command Modes

IPv4 address family configuration
 IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.3.0 | Support was added for the <i>k6</i> constant. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **metric weights** command to alter the default behavior of EIGRP routing and metric computation and allow the tuning of the EIGRP metric calculation for a particular ToS.

If *k5* equals 0, the composite EIGRP metric is computed according to the following formula:

$$\text{metric} = [k1 * \text{bandwidth} + (k2 * \text{bandwidth}) / (256 - \text{load}) + k3 * \text{delay}]$$

If $k5$ does not equal zero, an additional operation is performed:

$$\text{metric} = \text{metric} * [k5 / (\text{reliability} + k4)]$$

Bandwidth is inverse minimum bandwidth of the path in BPS scaled by a factor of $2.56 * 1012$. The range is from a 1200-bps line to 10 terabits per second.

Delay is in units of 10 microseconds. The range of delay is from 10 microseconds to 168 seconds. A delay of all 1s indicates that the network is unreachable.

The delay parameter is stored in a 32-bit field, in increments of 39.1 nanoseconds. The range of delay is from 1 (39.1 nanoseconds) to hexadecimal FFFFFFFF (decimal 4,294,967,040 nanoseconds). A delay of all 1s (that is, a delay of hexadecimal FFFFFFFF) indicates that the network is unreachable.

This table lists the default values used for several common media.

This command is supported only in 64 bit mode. The constraint $k6$ does not have effect in 64 bit mode.

Table 48: Bandwidth Values by Media Type

| Media Type | Delay | Bandwidth |
|------------|-----------------------------|----------------------|
| Satellite | 5120 (2 seconds) | 5120 (500 megabits) |
| Ethernet | 25600 (1 milliseconds [ms]) | 256000 (10 megabits) |
| 1.544 Mbps | 512000 (20,000 ms) | 1,657,856 bits |
| 64 kbps | 512000 (20,000 ms) | 40,000,000 bits |
| 56 kbps | 512000 (20,000 ms) | 45,714,176 bits |
| 10 kbps | 512000 (20,000 ms) | 256,000,000 bits |
| 1 kbps | 512000 (20,000 ms) | 2,560,000,000 bits |

Reliability is given as a fraction of 255. That is, 255 is a reliability of 100 percent or a perfectly stable link. Load is given as a fraction of 255. A load of 255 indicates a completely saturated link.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set the metric weights to change the default values:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# metric weights 0 2 0 2 0 0
```

This example shows how to set *k6* to the non-default value of *1*:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# metric weights 0 1 0 1 0 0 1
```

Related Commands

| Command | Description |
|--|--|
| metric maximum-hops, on page 639 | Advertises as unreachable those EIGRP VPN routes with a hop count higher than is specified by the command. |

neighbor (EIGRP)

To define a neighboring router with which to exchange Enhanced Interior Gateway Routing Protocol (EIGRP) information, use the **neighbor** command in interface configuration mode. To remove an entry, use the **no** form of this command.

neighbor *ip-address*

no neighbor *ip-address*

Syntax Description

| | |
|-------------------|--|
| <i>ip-address</i> | IP address of a peer router with which routing information is exchanged. |
|-------------------|--|

Command Default

No neighboring routers are defined.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **neighbor** command to permit the point-to-point (nonbroadcast) exchange of routing information.

If a neighbor is configured on an interface using the **neighbor** command, the interface stops sending or receiving multicast hello messages. However, the interface can send or receive unicast hello messages. So each neighbor on a LAN must be configured individually. Multiple **neighbor** commands can be used to specify additional neighbors or peers.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

This example shows how to permit the sending of EIGRP updates to specific neighbors. One copy of the update is generated for each neighbor:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv6
```

```
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# neighbor 172.20.1.2
```

Related Commands

| Command | Description |
|--|---|
| passive-interface (EIGRP), on page 654 | Disables sending and receiving "hello" messages on (EIGRP) interface. |

neighbor maximum-prefix

To limit the number of prefixes that are accepted from a single Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor or all EIGRP VPN neighbors, use the **neighbor maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

Single-Neighbor Configuration CLI

neighbor *ip-address* **maximum-prefix** *maximum* [*threshold*] [warning-only]

no neighbor *ip-address* **maximum-prefix**

All-Neighbor Configuration CLI

neighbor maximum-prefix *maximum* [*threshold*] [dampened] [reset-time *minutes*] [restart *minutes*] [restart-count *number*] [warning-only]

no neighbor maximum-prefix

Syntax Description

| | |
|------------------------------------|---|
| <i>ip-address</i> | (Optional) IP address of a single peer. |
| <i>maximum</i> | Maximum number of prefixes accepted. The range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router. |
| <i>threshold</i> | (Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent. |
| dampened | (Optional) Configures a decay penalty to be applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes. This keyword is disabled by default. |
| reset-time <i>minutes</i> | (Optional) Configures the router to reset the restart count to 0 after the default or configured reset time period has expired. The range is from 1 to 65535 minutes. |
| restart <i>minutes</i> | (Optional) Configures a time period in which the router does not form adjacencies or accept redistributed routes from the RIB after the maximum prefix limit has been exceeded. The range is from 1 to 65535 minutes. |
| restart-count <i>number</i> | (Optional) Configures the number of times a peering session can be automatically reestablished after the peering session has been torn down or a redistribute route has been cleared and relearned because the maximum prefix limit has been exceeded. Caution After the restart count threshold is crossed, you need to use the clear eigrp neighbors command to reestablish normal peering, redistribution, or both. |
| warning-only | (Optional) Configures the router to generate syslog messages only when the maximum prefix limit is reached, instead of terminating the peering session. |

Command Default

threshold: 75 percent
dampened : disabled
warning-only : disabled
reset-time : 15 minutes
restart : 5 minutes
restart-count : 3

Command Modes IPv4 VRF address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **neighbor maximum-prefix** command is configured to protect an individual peering session or all peering sessions. When this feature is enabled and the maximum prefix limit is exceeded, the router tears down the peering session, clears all routes that are learned from the peer, and then places the peer in a penalty state for the default or user-defined time period. After the penalty time period expires, normal peering is reestablished.



Note In EIGRP, **neighbor** commands have been used traditionally to configure static neighbors. In the context of this feature, however, the **neighbor maximum-prefix** command can be used to configure the maximum prefix limit for both statically configured and dynamically discovered neighbors.

When configuring the **neighbor maximum-prefix** command to protect a single peering session, just the maximum prefix limit, percentage threshold, and warning only configuration options can be configured. Session dampening, restart, and reset timers are configured on a global basis.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples

The following example shows how to configure the maximum prefix limit for a single peer. The maximum limit is set to 1000 prefixes, and the warning threshold is set to 80 percent. When the maximum prefix limit

is exceeded, the session with this peer is torn down, all routes learned from this peer are removed from the topology and routing tables, and this peer is placed in a penalty state for 5 minutes (default penalty value).

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# neighbor 10.0.0.1 maximum-prefix 1000 80
```

The following example shows how to configure the maximum prefix limit for all peers. The maximum limit is set to 10,000 prefixes, the warning threshold is set to 90 percent, the restart timer is set to 4 minutes, a decay penalty is configured for the restart timer with the **dampened** keyword, and all timers are configured to be reset to 0 every 60 minutes. When the maximum prefix limit is exceeded, all peering sessions are torn down, all routes learned from all peers are removed from the topology and routing tables, and all peers are placed in a penalty state for 4 minutes (user-defined penalty value). A dampening exponential decay penalty is also applied.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vrf1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# neighbor maximum-prefix 10000 90 dampened
reset-time 60 restart4
```

Related Commands

| Command | Description |
|---|---|
| clear eigrp neighbors , on page 615 | Deletes EIGRP and VRF neighbor entries from the appropriate tables. |

next-hop-self disable

To instruct the Enhanced Interior Gateway Routing Protocol (EIGRP) process to use the received next-hop value when advertising the routes, use the **next-hop-self disable** command in interface configuration mode. To revert to the default, use the **no** form of this command.

next-hop-self disable

no next-hop-self disable

Syntax Description This command has no keywords or arguments.

Command Default EIGRP always sets the IP next-hop value to be itself.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

EIGRP, by default, sets the IP next-hop value to be itself for routes that it is advertising, even when advertising those routes on the same interface from which learned them. To change this default, you must use the **next-hop-self disable** interface configuration command to instruct EIGRP to use the received next-hop value when advertising these routes.

The **next-hop-self disable** feature is not available for redistributed routes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples The following example shows how to change the default IP next-hop value and instruct EIGRP to use the received next-hop value:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af) # interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if) # next-hop-self disable
```

nsf disable (EIGRP)

To disable Enhanced Interior Gateway Routing Protocol (EIGRP) nonstop forwarding (NSF), use the **nsf disable** command in appropriate configuration mode. To re-enable the EIGRP NSF from a disabled state, use the **no** form of this command.

nsf disable

no nsf disable

Syntax Description This command has no keywords or arguments.

Command Default NSF in EIGRP is enabled.

Command Modes

- Router Configuration
- IPv4 address family configuration
- IPv6 address family configuration
- IPv4 VRF address family configuration
- IPv6 VRF address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.1.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If NSF is to be disabled for both IPv4 and IPv6 address families of all VRFs, use the **nsf disable** command at router configuration mode.

If NSF is to be disabled for a specific address family under a specific VRF, use the **nsf disable** command at address family configuration mode.

If NSF is disabled, EIGRP will not be able to retain the routes learnt from its peers and may result in traffic loss during ISSU.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | eigrp | read, write |

Examples

This example shows how to disable NSF for all address families under all VRF's:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)#nsf disable
```

This example shows how to disable NSF for IPv4 address family of VRF v1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)#vrf v1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)#address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)#nsf disable
```

Related Commands

| Command | Description |
|--|---|
| router eigrp , on page 662 | Configures a routing process and enters router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP). |

passive-interface (EIGRP)

To disable sending and receiving "hello" messages on an Enhanced Interior Gateway Routing Protocol (EIGRP) interface and to disable formation of neighbors on the interface, use the **passive-interface** command in interface configuration mode. To reenble sending and receiving "hello messages", use the **no** form of this command.

passive-interface
no passive-interface

Syntax Description This command has no keywords or arguments.

Command Default **passive-interface** command is disabled on an interface.

Command Modes Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **passive-interface** command to disable sending of "hello" messages. The particular subnet on that interface will continue to be advertised by EIGRP to neighbors on other interfaces.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

This example shows how to configure **passive-interface** command on GigabitEthernet interface 0/6/5/0:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv6
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/6/5/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# passive-interface
```

redistribute (EIGRP)

To inject routes from one routing domain into the Enhanced Interior Gateway Routing Protocol (EIGRP), use the **redistribute** command in the appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

```
redistribute {{bgp|connected|isis|ospf|rip|static|eigrp}} [as-number|instance-name][ route-policy name]
```

```
no redistribute
```

Syntax Description

| | |
|---|--|
| bgp | Distributes routes from the BGP protocol. |
| connected | Distributes routes that are established automatically by virtue of having enabled IP on an interface. |
| isis | Distributes routes from the IS-IS protocol. |
| ospf | Distributes routes from the OSPF protocol. This protocol is supported in the IPv4 address family. |
| static | Redistributes IP static routes. |
| eigrp | Redistributes routes from other EIGRP autonomous-systems. |
| <i>as-number</i> <i>instance-name</i> | Represents one of the following three options: For the bgp keyword: Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. For the isis keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the ospf keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. For the eigrp keyword, a 2-byte autonomous-system number. Range is 1 to 65535. |
| route-policy <i>name</i> | (Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to EIGRP. |

Command Default Route redistribution is disabled.

Command Modes IPv4 address family configuration
IPv4 VRF address family configuration

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |
| Release 5.1 | The eigrp keyword was added. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Redistributed routing information should always be filtered by the **route- policy name** keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed by EIGRP.

A default metric is usually required to redistribute routes from another protocol into EIGRP. The metric is configured through the **default-metric** command or under the route policy configured with the **redistribute** command. The two exceptions to this requirement are when EIGRP redistributes BGP routes on a provider edge (PE) router in an MPLS-VPN scenario and when EIGRP redistributes EIGRP routes from another autonomous-system number. In case of MPLS-VPN scenario, if the originating protocol of the route is EIGRP with the same autonomous-system, the metric would be learned automatically from the extended communities of the BGP route. In case of redistribution from EIGRP belonging to different autonomous-system, the metric would be learned automatically from the extended communities from RIB.

For information about routing policies, see the *Routing Policy Commands on Cisco ASR 9000 Series Router* module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples This example shows how to cause BGP routes to be redistributed into an EIGRP autonomous system:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# redistribute bgp 100
```

This example shows how to redistribute the specified IS-IS process routes into an EIGRP autonomous system within a VPN routing and forwarding instance (VRF). The IS-IS routes are redistributed using route policy 3.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 109
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute isis 108 route-policy 3
```

This example shows how to cause EIGRP routes from autonomous-system 100 to be redistributed into EIGRP autonomous-system 1.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# redistribute eigrp 100
```

Related Commands

| Command | Description |
|---|----------------------------|
| default-metric (EIGRP), on page 621 | Sets metrics for an EIGRP. |

redistribute maximum-prefix

To limit the number of prefixes redistributed into an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **redistribute maximum-prefix** command in IPv4 VRF address family configuration mode. To disable this function, use the **no** form of this command.

redistribute maximum-prefix *maximum* [*threshold*] [[**dampened**] [**reset-time** *minutes*] [**restart** *minutes*] [**restart-count** *number*]] [**warning-only**]]

no redistribute maximum-prefix

Syntax Description

| | |
|------------------------------------|---|
| maximum | Maximum number of prefixes that are redistributed into EIGRP under an address-family. The range is from 1 to 4294967295. The number of prefixes that can be configured is limited only by the available system resources on the router. |
| threshold | (Optional) Syslog warning messages are specified as a percentage of the maximum prefix limit that was exceeded. The prefix percentage number range is from 1 to 100. The default is 75 percent. |
| restart <i>minutes</i> | (Optional) Configures a time period in which the router will not form adjacencies or accept redistributed routes from the RIB after the <i>maximum-prefix</i> limit has been exceeded. The value for the minutes argument is from 1 to 65535 minutes. |
| restart-count <i>number</i> | (Optional) Configures the number of times a peering session can be automatically reestablished after the peering session has been torn down or after the redistribute route has been cleared and relearned because the maximum prefix limit has been exceeded. After the restart count threshold has been crossed, you will need to enter the process restart eigrp command to reestablish normal peering, redistribution, or both. |
| reset-time <i>minutes</i> | (Optional) Configures the router to reset the restart count to 0 after the default or configured reset time period has expired. The value for the minutes argument is from 1 to 65535 minutes. |
| dampened | (Optional) Configures a decay penalty to be applied to the restart time period each time the maximum prefix limit is exceeded. The half-life for the decay penalty is 150 percent of the default or user-defined restart time value in minutes. |
| warning-only | (Optional) Configures the router to only generate syslog messages when the maximum prefix limit is reached, instead of suspending redistribution. |

Command Default

threshold: 75 percent
warning-only : disabled
reset-time : 15 minutes

restart : 5 minutes
restart-count : 3
dampened: disabled

Command Modes IPv4 VRF address family configuration

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **redistribute maximum-prefix** command to configure limit prefixes learned through redistribution. When the maximum prefix limit is exceeded, all routes learned from the Routing Information Base (RIB) are discarded and redistribution is suspended for the default or user-defined time period. The maximum prefix limit that can be configured for redistributed prefixes is limited only by the available system resources on the router.

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples The following example shows how to configure the maximum prefix limit for routes learned through redistribution. The maximum limit is set to 5000 prefixes, and the warning threshold is set to 95 percent. When the number of prefixes learned through redistribution reaches 4750 (95 percent of 5000), warning messages are displayed in the console. Because the **warning-only** keyword is configured, the topology and routing tables are not cleared and route redistribution is not placed in a penalty state.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 100
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# redistribute maximum-prefix 5000 95 warning-only
```

| Command | Description |
|---------|--|
| process | To start, terminate, or restart a process. |

route-policy (EIGRP)

To apply a routing policy to updates advertised to or received from an Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor, use the **route-policy** command in the appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *route-policy-name* {**in**|**out**}

no route-policy *route-policy-name* {**in**|**out**}

Syntax Description

| | |
|--------------------------|------------------------------------|
| <i>route-policy-name</i> | Name of route policy. |
| in | Applies policy to inbound routes. |
| out | Applies policy to outbound routes. |

Command Default

No policy is applied.

Command Modes

IPv4 address family configuration
 IPv4 VRF address family configuration
 Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to apply the IN-IPv4 policy to inbound IP Version 4 (IPv4) routes:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af) # route-policy IN-IPv4 in
```

router eigrp

To configure a routing process and enter router configuration mode for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **router eigrp** command in global configuration mode. To turn off the EIGRP routing process, use the **no** form of this command.

router eigrp { *instance-autonomous-system-number* | *virtual-instance-name* }

no router eigrp { *instance-autonomous-system-number* | *virtual-instance-name* }

Syntax Description

instance-autonomous-system-number EIGRP instance autonomous-system number. This is used as the autonomous-system number for the default/global VRF. Valid range is 1 to 65535.

virtual-instance-name EIGRP virtual instance name. The name can be up to 32 characters long and cannot be a number greater than 65535.

Command Default

No routing process is defined.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 5.1 | The virtual-instance-name variable was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

- An explicit autonomous-system configuration is required for the default/global VRF when EIGRP is configured using virtual instance-name.
- Up to 4 EIGRP instances can be configured.
-

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example configures EIGRP routing process 109 (109 is the autonomous-system number of default/global VRF):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 109  
RP/0/RSP0/CPU0:router(config-eigrp)#
```

This example configures an EIGRP routing process and assigns it the name "name":

```
RP/0/RSP0/CPU0:router(config)# router eigrp name  
RP/0/RSP0/CPU0:router(config-eigrp)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af)# autonomous-system 109
```

router-id (EIGRP)

To configure a router ID for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **router-id** command in the appropriate configuration mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command.

router-id *router-id*

no router-id

Syntax Description

| | |
|------------------|---|
| <i>router-id</i> | 32-bit router ID value specified in four-part, dotted-decimal notation. |
|------------------|---|

Command Default

If this command is not configured, EIGRP chooses an IPv4 address as the router ID from one of its interfaces.

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

It is good practice to use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that EIGRP can function regardless of the interface address configuration.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to assign the IP address of 172.20.1.1 to the EIGRP process 1:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# router-id 172.20.1.1
```

show eigrp accounting

To display prefix accounting information for Enhanced Interior Gateway Routing Protocol (EIGRP) processes, use the **show eigrp accounting** command in EXEC mode.

show eigrp [*as-number*] [**vrf** {*vrf-name*| **all**}] [**ipv4**| **ipv6**] **accounting**

Syntax Description

| | |
|---|--|
| <i>as-number</i> | (Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535. |
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| [ipv4] | (Optional) Specifies the IPv4 address family. |

Command Default

This command has no arguments or keywords.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| EIGRP | read |

Examples

The following is sample output from the **show eigrp accounting** command:

```
RP/0/RSP0/CPU0:router# show eigrp accounting
IP-EIGRP accounting for AS(100)/ID(10.0.2.1) Routing Table: RED
Total Prefix Count: 4 States: A-Adjacency, P-Pending, D-Down
State Address/Source Interface Prefix Restart Restart/
```

```

Count      Count      Reset (s)
P   Redistributed   ----      0           3           211
A   10.0.1.2        Gi0/6/0/0/  2           0            84
P   10.0.2.4        Gi00/2/0/3  0           2           114
D   10.0.1.3        Gi0/6/0/0   0           3            0

```



Note Connected and summary routes are not listed individually in the output of this command but are counted in the total aggregate count for each process.

This table describes the significant fields shown in the display.

Table 49: show eigrp accounting Field Descriptions

| Field | Description |
|--|---|
| EIGRP accounting for AS | Identifies the EIGRP instance along with the AS number, router ID and table ID. |
| Total Prefix Count | Shows the aggregate sum of the prefixes in an EIGRP instance topology table. The count includes prefixes learned from all neighbors or from redistribution. |
| States: A-Adjacency, P-Pending, D-Down | <p>A-Adjacency: Indicates a stable adjacency with the neighbor or a normal redistribution state.</p> <p>P-Pending: Neighbor adjacency or redistribution is suspended or in a penalized state because the maximum prefix limit was exceeded.</p> <p>D-Down: Neighbor adjacency or redistribution is suspended permanently until a manual reset is performed with the clear route command.</p> |
| Address/Source | Shows the peer IP address of the redistribution source. |
| Prefix Count | <p>Displays the total number of learned prefixes by source.</p> <p>Note Routes can be learned for the same prefix from multiple sources, and the sum of all prefix counts in this column may be greater than the figure displayed in the "Prefix Count" field.</p> |
| Restart Count | Number of times a route source exceeded the maximum prefix limit. |
| Restart Reset(s) | Displays the time, in seconds, that a route source is in a P (penalized) state. If the route source is in an A (stable or normal) state, the displayed time, in seconds, is the time period until penalization history is reset. |

show eigrp interfaces

To display information about interfaces configured for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **show eigrp interfaces** command in EXEC mode.

show eigrp [*as-number*] [**vrf**{*vrf-name*| **all**}] [**ipv4**| **ipv6**] **interfaces** [*type interface-path-id*] [**detail**]

Syntax Description

| | |
|---|---|
| as-number | (Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| [ipv4] | (Optional) Specifies the IPv4 address family. |
| type | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| interface-path-id | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| detail | (Optional) Displays detailed EIGRP interface information. |

Command Default

This command has no arguments or keywords.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|---|
| Release 3.9.0 | As plain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.3.0 | The command output was modified to display relevant wide metric information. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show eigrp interfaces** command to determine on which interfaces EIGRP is active and learn information about EIGRP related to those interfaces.

If an interface is specified, only that interface is displayed. Otherwise, all interfaces on which EIGRP is running are displayed.

If an autonomous system is specified, only the routing process for the specified autonomous system is displayed. Otherwise, all EIGRP processes are displayed.

Task ID

| Task ID | Operations |
|---------|------------|
| EIGRP | read |

Examples

The following is sample output from the **show eigrp interfaces** command:

```
RP/0/RSP0/CPU0:router# show eigrp interfaces
IP EIGRP interfaces for process 1

Interface          Peers    Xmit Queue   Mean   Pacing Time   Multicast   Pending
                  Un/Reliable SRTT    Un/Reliable   Flow Timer  Routes
-----
Gi0/6/0/2.212      0         0/0          0      11/434         0           0
Gi0/6/0/0          1         0/0          337    0/10           0           0
  Gi0/2/0/3        1         0/0          10     1/63           103         0
Gi0/6/2/5          1         0/0          330    0/16           0           0
```

This table describes the significant fields shown in the display.

Table 50: show eigrp interfaces Field Descriptions

| Field | Description |
|-----------|---|
| Interface | Interface over which EIGRP is configured. |

| Field | Description |
|-------------------------|--|
| Peers | Number of directly connected EIGRP neighbors. |
| Xmit Queue Un/Reliable | Number of packets remaining in the unreliable and reliable transmit queues. |
| Mean SRTT | Mean smoothed round-trip time (SRTT) internal (in milliseconds). |
| Pacing Time Un/Reliable | Pacing time used to determine when EIGRP packets should be sent out the interface (unreliable and reliable packets). |
| Multicast Flow Timer | Maximum number of seconds in which the router sends multicast EIGRP packets. |
| Pending Routes | Number of routes in the packets in the transmit queue waiting to be sent. |

The following is sample output from the **show eigrp interfaces** command when issued with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show eigrp interfaces detail
IPv4-EIGRP interfaces for AS(100)

Interface          Peers  Xmit Queue  Mean   Pacing Time  Multicast   Pending
                   Un/Reliable SRTT    Un/Reliable  Flow Timer   Routes
-----          -
Lo0                0      0/0         0      640/640      0           0
  Hello interval is 5 sec, hold time is 15 sec
  Next xmit serial <none>
  Un/reliable mcasts: 0/0  Un/reliable ucasts: 0/0
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 0  Out-of-sequence rcvd: 0
  Bandwidth percent is 50
  Total packets received: 0
  Authentication mode: MD5  Key chain: key1
  No active key found in keychain database
  Valid authenticated packets received: 0
  Packets dropped due to wrong keychain config: 0
  Packets dropped due to missing authentication: 0
  Packets dropped due to invalid authentication: 0
  Effective Metric:
    Bandwidth: 10000000, Delay: 500, Reliability: 255, Load: 1, MTU: 1514
```

This table describes the significant fields shown in the display.

Table 51: show eigrp interfaces detail Field Descriptions

| Field | Description |
|----------------|-------------------------------------|
| Hello interval | Hello packet transmission interval. |

| Field | Description |
|---|--|
| hold time | Hold time announced to neighbors. If neighbors do not get hello packets from the router for this period, neighbors declare that the neighbor relationship is down. |
| Next xmit serial | Next transmission serial number. |
| Un/reliable mcasts | Number of unreliable and reliable multicast packets sent on this interface. |
| Un/reliable ucasts | Number of unreliable and reliable unicast packets sent on this interface. |
| Mcast exceptions | Number of multicast exceptions (sequence TLVs). |
| CR packets | Number of packets sent with the conditional receive bit set. |
| ACKs suppresses | Number of ACK packets suppressed. |
| Retransmissions | Number of retransmissions sent on this interface. |
| Out-of-sequence rcvd | Number of packets received out of sequence. |
| Bandwidth percent | Configured percent of bandwidth. |
| Authentication | Mode of authentication. |
| Valid authenticated packets received | Number of valid authentication packets. |
| Packets dropped due to wrong keychain config | Number of packets dropped due to wrong keychain configuration. |
| Packets dropped due to missing authentication | Number of packets dropped due to missing authentication. |
| Packets dropped due to invalid authentication | Number of packets dropped due to invalid authentication. |

This is sample **show eigrp interfaces** command output to display relevant wide metric information:

```
RP/0/RSP0/CPU0:router#show eigrp interfaces gigabitEthernet 0/0/0/0 detail
IPv4-EIGRP interfaces for AS(1)
Interface          Peers   Xmit Queue   Mean   Pacing Time   Multicast   Pending
                   Un/Reliable SRTT   Un/Reliable  Flow Timer   Routes
Gi0/0/0/0          1       0/0         420    0/10         2096        0
Hello interval is 5 sec, hold time is 15 sec
Next xmit serial <none>
Un/reliable mcasts: 0/6  Un/reliable ucasts: 4/2
```

show eigrp interfaces

```
Mcast exceptions: 0 CR packets: 0 ACKs suppressed: 1
Retransmissions sent: 1 Out-of-sequence rcvd: 0
Bandwidth percent is 50
Total packets received: 1563
Authentication mode is not set
Classic peers: 0 Wide peers: 1
Effective Metric:
  Bandwidth: 1000000 Kbit, Delay: 10000000 picosecond,
  Reliability: 255, Load: 1, MTU: 1500
```

Related Commands

| Command | Description |
|---|---|
| show eigrp neighbors, on page 673 | Displays the neighbors discovered by EIGRP. |

show eigrp neighbors

To display information about neighbors discovered by Enhanced Interior Gateway Routing Protocol (EIGRP), use the **show eigrp neighbors** command in EXEC mode.

```
show eigrp as-number vrf {vrf-name | all} ipv4 ipv6
```

Syntax Description

| | |
|---|---|
| <i>as-number</i> | (Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535. |
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| [ipv4] | (Optional) Specifies the IPv4 address family. |
| detail | (Optional) Displays detailed EIGRP neighbor information. |
| type | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| static | (Optional) Displays static routes. |

Command Default No default behavior or values

Command Modes EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.3.0 | The command output was modified to display relevant wide metric information. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show eigrp neighbors** command to determine when neighbors become active and inactive. This command is also useful for debugging certain types of transport problems.

Task ID

| Task ID | Operations |
|---------|------------|
| EIGRP | read |

Examples

The following is sample output from the **show eigrp neighbors** command:

```
RP/0/RSP0/CPU0:router# show eigrp neighbors
IP-EIGRP Neighbors for process 77
Address                Interface    Holdtime  Uptime    Q      Seq  SRTT  RTO
                   (secs)    (h:m:s)  Count    Num    (ms)  (ms)
172.16.81.28          Gi0/2/0/3   13        0:00:41   0       11    4     20
172.16.80.28          Gi0/6/0/0   14        0:02:01   0       10   12    24
172.16.80.31          Gi0/6/2/5   12        0:02:02   0        4    5     20
RP/0/RSP0/CPU0:router#
```

This table describes the significant fields shown in the display.

Table 52: show eigrp neighbors Field Descriptions

| Field | Description |
|-----------|--|
| process | Autonomous system number specified in the router configuration command. |
| Address | IP address of the EIGRP peer. |
| Interface | Interface on which the router is receiving hello packets from the peer. |
| Holdtime | Length of time (in seconds) that the Cisco IOS XR software waits to hear from the peer before declaring that the peer is down. |
| Uptime | Elapsed time (in hours, minutes, and seconds) since the local router first heard from this neighbor. |

| Field | Description |
|---------|---|
| Q Count | Number of EIGRP packets (update, query, and reply) that the software waits to send. |
| Seq Num | Sequence number of the last update, query, or reply packet that was received from this neighbor. |
| SRTT | Smoothed round-trip time. This is the number of milliseconds required for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet. |
| RTO | Retransmission timeout (in milliseconds). This is the amount of time the software waits before resending a packet from the retransmission queue to a neighbor. |

The following is sample output from the **show eigrp neighbors** command when issued with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show eigrp neighbors detail
IP-EIGRP neighbors for AS 1
H   Address                Interface      Hold Uptime   SRTT  RTO  Q  Seq
   (sec)                    (ms)          (sec)         (ms)  Cnt  Num
0   11.0.0.10                Gi0/6/0/0     14 01:00:52   3     200  0  10
Version 12.4/1.2, Retrans: 0, Retries: 0, Prefixes: 3
```

This table describes the significant fields shown in the display.

Table 53: show eigrp neighbors detail Field Descriptions

| Field | Description |
|----------|---|
| Version | Version of EIGRP software (major.minor) running on the node and neighbor. |
| Retrans | Number of retransmissions sent to this neighbor. |
| Retries | Number of retransmissions sent to this neighbor since the last acknowledgement (ACK). |
| Prefixes | Number of prefixes learned from this neighbor. |

This is sample output from **show eigrp neighbors** command to display relevant wide metric information:

```
RP/0/RSP0/CPU0:router#show eigrp neighbors detail
Mon Oct 31 21:23:37.996 IST
IPv4-EIGRP neighbors for AS(1) vrf default
```

show eigrp neighbors

| H | Address | Interface | Hold Uptime | SRTT | RTO | Q | Seq |
|---|-------------|-----------|-------------|------|------|-----|-----|
| | | | (sec) | (ms) | | Cnt | Num |
| 0 | 10.10.10.11 | Gi0/0/0/0 | 12 01:20:40 | 420 | 2520 | 0 | 5 |

Version 3.3/2.0, Retrans: 1, Retries: 0, Prefixes: 2

show eigrp topology

To display the Enhanced Interior Gateway Routing Protocol (EIGRP) topology table, use the **show eigrp topology** command in EXEC mode.

```
show eigrp [as-number] [vrf{vrf-name| all}] [ipv4| ipv6] topology [ip-address mask] {active| all-links| detail-links| pending| summary| zero-successors}
```

Syntax Description

| | |
|---|---|
| <i>as-number</i> | (Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| [ipv4] | (Optional) Specifies the IPv4 address family. |
| <i>ip-address</i> | (Optional) IP address in four-part, dotted-decimal notation. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |
| active | (Optional) Displays only active entries in the EIGRP topology table. |
| all-links | (Optional) Displays all entries in the EIGRP topology table. |
| detail-links | (Optional) Displays detailed information for all entries in the EIGRP topology table. |
| pending | (Optional) Displays all entries in the EIGRP topology table that are waiting for an update from a neighbor or are waiting to reply to a neighbor. |
| summary | (Optional) Displays a summary of the EIGRP topology table. |
| zero-successors | (Optional) Displays available routes in the EIGRP topology table. |

Command Modes EXEC

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.3.0 | The command output was modified to display relevant wide metric information. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the **show eigrp topology** command is used without any keywords or arguments, only routes that are feasible successors are displayed.

The **show eigrp topology** command can be used to determine Diffusing Update Algorithm (DUAL) states and to debug possible DUAL problems.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | EIGRP | read |

Examples The following is sample output from the **show eigrp topology** command. The EIGRP metrics for specified internal routes and external routes are displayed.

```
RP/0/RSP0/CPU0:router# show eigrp topology 10.2.1.0/24
IP-EIGRP (AS 1): Topology entry for 10.2.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 281600
  Routing Descriptor Blocks:
  0.0.0.0 (GigabitEthernet0/6/0/0), from Connected, Send flag is 0x0
    Composite metric is (281600/0), Route is Internal !This is the internal route.
    Vector metric:
      Minimum bandwidth is 10000 Kbit
      Total delay is 1000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 0

RP/0/RSP0/CPU0:router# show eigrp topology 10.4.80.0/20
IP-EIGRP (AS 1): Topology entry for 10.4.80.0/20
```

```

State is Passive, Query origin flag is 1, 1 Successor(s), FD is 409600
Routing Descriptor Blocks:
10.2.1.1 (GigabitEthernet0/6/0/0), from 10.2.1.1, Send flag is 0x0
  Composite metric is (409600/128256), Route is External
  Vector metric:
    Minimum bandwidth is 10000 Kbit
    Total delay is 6000 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 1
  External data:
    Originating router is 10.89.245.1
    AS number of route is 0
    External protocol is Connected, external metric is 0
    Administrator tag is 0 (0x00000000)

```

This table describes the significant fields shown in the display.

Table 54: show eigrp topology Field Descriptions

| Field | Description |
|-------------------------------------|--|
| Query origin | Query origin state. |
| Successors | Number of feasible successors for this prefix. |
| FD | Feasible distance for this prefix. |
| 10.2.1.1 (Gi0/0) | Next hop and interface from which this path was learned. |
| from 10.2.1.1 | Information source for this path. |
| Send flag | Indicates if the sending of this prefix is pending to this neighbor. |
| Composite Metric (409600/128256) | The first number is the EIGRP metric that represents the cost to the destination. The second number is the EIGRP metric that this peer advertised. |
| Route is | Type of route (internal or external). |
| Vector Metric | Shows the metric (bandwidth, delay, reliability, load, MTU, and hop count) advertised by the neighbor. |
| External Data | Shows the external information (originating router ID, AS number, external protocol, metric, and tag) advertised by the neighbor. |

This is sample **show eigrp topology** command output to display relevant wide metric information:

```

RP/0/RSP0/CPU0:router#show eigrp topology 1.1.1.0/24
IPv4-EIGRP AS(1): Topology entry for 1.1.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 82329600, RIB is 643200
  Routing Descriptor Blocks:

```

```
show eigrp topology
```

```
10.10.10.11 (Ethernet0/0), from 10.10.10.11, Send flag is 0x0  
Composite metric is (82329600/163840), Route is Internal  
Vector metric:  
  Minimum bandwidth is 16000 Kbit  
  Total delay is 631250000 picosecond  
  Reliability is 255/255  
  Load is 1/255  
  Minimum MTU is 1500  
  Hop count is 1  
  Originating router is 1.1.1.1
```

show eigrp traffic

To display the number of Enhanced Interior Gateway Routing Protocol (EIGRP) packets sent and received, use the **show eigrp traffic** command in EXEC mode.

```
show eigrp [as-number] [vrf{vrf-name| all}][ipv4| ipv6] traffic
```

Syntax Description

| | |
|---|--|
| <i>as-number</i> | (Optional) Autonomous system number. This option is available when a VPN routing and forwarding (VRF) instance is not specified. Range is from 1 to 65535. |
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding instance (VRF) or all VRF instances. |
| [ipv4] | (Optional) Specifies the IPv4 address family. |

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show eigrp traffic** command to find the number of packets sent and received.

In addition, this command is useful in determining whether packets from one node are not reaching the neighboring node due to connectivity or configuration problems.

Task ID

| Task ID | Operations |
|---------|------------|
| EIGRP | read |

Examples

The following is sample output from the **show eigrp traffic** command:

```
RP/0/RSP0/CPU0:router# show eigrp traffic
IP-EIGRP Traffic Statistics for AS 1
```

```

Hellos sent/received: 736/797
Updates sent/received: 6/6
Queries sent/received: 0/1
Replies sent/received: 1/0
Acks sent/received: 6/6
Input queue high water mark 0, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0

```

This table describes the significant fields shown in the display.

Table 55: show eigrp traffic Field Descriptions

| Field | Description |
|-----------------------------|--|
| AS | Autonomous system number specified in the router eigrp command. |
| Hellos sent/received | Number of hello packets sent and received. |
| Updates sent/received | Number of update packets sent and received. |
| Queries sent/received | Number of query packets sent and received. |
| Replies sent/received | Number of reply packets sent and received. |
| Acks sent/received | Number of acknowledgment packets sent and received. |
| Input queue high water mark | Maximum number of packets in the input queue and number of drops. |
| SIA-Queries sent/received | Number of Stuck-in-Active query packets sent and received. |
| SIA-Replies sent/received | Number of Stuck-in-Active reply packets sent and received. |

show protocols (EIGRP)

To display information about the Enhanced Interior Gateway Routing Protocol (EIGRP) process configuration, use the **show protocols** command in EXEC mode.

```
show protocols [ipv4 | afi-all] [all | protocol] [default-context | vrf vrf-name] [private]
```

Syntax Description

| | |
|------------------------|--|
| ipv4 | (Optional) Specifies an IPv4 address family. |
| afi-all | (Optional) Specifies all address families. |
| all | (Optional) Specifies all protocols for a given address family. |
| protocol | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are eigrp , bgp , isis , ospf , and rip . |
| default-context | (Optional) Displays default context information. This keyword is available when the eigrp or rip protocol is specified. |
| vrf vrf-name | (Optional) Displays VPN routing and forwarding (VRF) information for the specified process. This keyword is available when the eigrp or rip protocol is specified. |
| private | (Optional) Displays private EIGRP data. This keyword is available when the eigrp is specified. |

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.3.0 | The command output was modified to display relevant wide metric information. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command is designed to summarize the important characteristics of the running protocol, and command output varies depending on the specific protocol selected.

For EIGRP, the command output lists the instance number, default AS context, router ID, default networks, distance, maximum paths, and so on.

Task ID

| Task ID | Operations |
|---------|------------|
| EIGRP | read |

Examples

The following is sample output from the **show protocols eigrp** command:

```
RP/0/RSP0/CPU0:router# show protocols eigrp

Routing Protocol: EIGRP, instance 1
Default context AS: 1, Router ID: 192.168.0.22
Address Family: IPv4
  Default networks not flagged in outgoing updates
  Default networks not accepted from incoming updates
  Distance: internal 90, external 170
  Maximum paths: 4
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  EIGRP NSF: enabled
  NSF-aware route hold timer is 240s
  NSF signal timer is 20s
  NSF converge timer is 120s
  Time since last restart is 01:01:21
  SIA Active timer is 180s
  Interfaces:
  GigabitEthernet0/6/0/0
```

This table describes the significant fields shown in the display.

Table 56: show protocols Field Descriptions

| Field | Descriptions |
|----------|---|
| instance | <p>AS number of the instance.</p> <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |

| Field | Descriptions |
|----------------------------|--|
| AS | AS number of this context. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| Address Family | Address family for which the configuration status is shown. |
| Default Networks Candidate | Default network acceptance and announcement behavior. |
| Distance | Administrative distance of EIGRP routes. |
| Maximum paths | Maximum paths installed in RIB for a route. |
| Metric Weight | Current metric weights used by EIGRP. |
| Maximum hopcount | Maximum hop count accepted by EIGRP. |
| Variance | Metric variance used to find feasible paths for a route. |
| Route hold time | Time duration for which routes learned from a neighbor are held without deletion while the neighbor is undergoing a graceful restart. |
| signal time | Nonstop forwarding signal time. |
| converge time | Nonstop forwarding convergence time. |
| SIA Active time | Active time period for SIA. |
| Interfaces | List of interfaces configured for EIGRP. |

This is sample output from the **show protocols eigrp** command with wide metric information:

```
RP/0/RSP0/CPU0:router#show protocols eigrp
Routing Protocol: EIGRP, instance 1
Default context AS: 1, Router ID: 3.3.3.3
Address Family: IPv4
  Default networks not flagged in outgoing updates
  Default networks not accepted from incoming updates
  Distance: internal 90, external 170
  Maximum paths: 4
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0, K6=0
Metric rib-scale 128
Metric version 64bit
```

show protocols (EIGRP)

```

EIGRP maximum hopcount 100
EIGRP maximum metric variance 1

EIGRP NSF: enabled
NSF-aware route hold timer is 480s
NSF signal timer is 20s
NSF converge timer is 300s
Time since last restart is 01:40:15
SIA Active timer is 180s
Interfaces:
  Loopback0
  GigabitEthernet0/0/0/0

```

This table describes the significant fields shown in the display.

| Field | Descriptions |
|----------------------|-----------------------------------|
| K6=0 | K6 constraint value |
| Metric rib-scale 128 | Metric rib-scale scale-down-value |
| Metric version 64bit | Metric version in bits |

site-of-origin (EIGRP)

To configure the Site of Origin (SoO) filtering on an Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **site-of-origin** command in interface configuration mode. To disable SoO filtering on an interface, use the **no** form of this command.

site-of-origin {*as-number* : *number*| *ip-address* : *number*}

no site-of-origin

Syntax Description

| | |
|---------------------|--|
| <i>as-number</i> : | Autonomous system number. Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. The colon is used to separate the autonomous system number and network number. |
| <i>number</i> | Network number. Range is from 0 to 4294967295 when a 2-byte AS number is used. Range is from 0 to 65535 when a 4-byte AS number is used. |
| <i>ip-address</i> : | IP address in four-part, dotted-decimal notation. The colon is used to separate the IP address and network number. |

Command Default

No default behavior or values

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An EIGRP process must be capable of retrieving the SoO attribute on routes redistributed from the Border Gateway Protocol (BGP) when required to support complex topologies that include MPLS VPN links between sites with backdoor links.

Use the **site-of-origin** command to set an SoO BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SoO extended community uniquely identifies the site from which a provider edge (PE) router has learned a route.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to configure SoO filtering on an EIGRP interface:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) vrf customer1
RP/0/RSP0/CPU0:router(config-eigrp-vrf) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp--vrf-af-if)# site-of-origin 10.0.0.1:20
```

split-horizon disable (EIGRP)

To disable split horizon for an Enhanced Interior Gateway Routing Protocol (EIGRP) process, use the **split-horizon disable** command in interface configuration mode. To enable split horizon, use the **no** form of this command.

split-horizon disable

no split-horizon disable

Syntax Description This command has no keywords or arguments.

Command Default Split horizon is enabled for an EIGRP process.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples The following example shows how to disable split horizon on a GigabitEthernet link:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# split-horizon disable
```

stub (EIGRP)

To configure a router as a stub for Enhanced Interior Gateway Routing Protocol (EIGRP), use the **stub** command in the appropriate configuration mode. To disable this function, use the **no** form of this command.

stub [receive-only] [connected] [redistributed] [static] [summary]

no stub [receive-only] [connected] [redistributed] [static] [summary]

Syntax Description

| | |
|----------------------|---|
| receive-only | (Optional) Sets the router as a receive-only neighbor. |
| connected | (Optional) Advertises connected routes. |
| redistributed | (Optional) Advertises redistributed routes from other protocols and autonomous systems. |
| static | (Optional) Advertises static routes. |
| summary | (Optional) Advertises summary routes. |

Command Default

Stub routing is disabled.

When stub routing is specified, connected and summary routes are advertised by default.

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **stub** command to configure a router as a stub in which the router directs all IP traffic to a distribution router.

The **stub** command can be modified with several options, and these options can be used in any combination except for the **receive-only** keyword.

The **receive-only** keyword restricts the router from sharing any of its routes with any other router in that EIGRP autonomous system and does not permit any other option to be specified because it prevents any type of route from being sent. The four other optional keywords (**connected** , **static** , **summary** , and **redistributed**

) can be used in any combination but cannot be used with the **receive-only** keyword. If any of these four keywords is used with the **stub** command, only the route types specified by the particular keyword or keywords are sent. Route types specified by the nonused keyword or keywords are not sent.

The **connected** keyword permits EIGRP stub routing to send connected routes. If all the connected routes are not covered by EIGRP interfaces, it may be necessary to redistribute connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default.

The **static** keyword permits EIGRP stub routing to send static routes. Without the configuration of this option, EIGRP does not send any static routes. You may still need to redistribute static routes with the **redistribute static** command.

The **summary** keyword permits EIGRP stub routing to send summary routes. Summary routes can be created manually with the **summary address** command or automatically at a major network border router with the **auto-summary** command enabled. This option is enabled by default.

The **redistributed** keyword permits EIGRP stub routing to send other routing protocols and autonomous systems. Without the configuration of this option, EIGRP does not advertise redistributed routes.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to configure, as a stub, the router that advertises connected and summary routes:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# stub
```

The following example shows how to configure the router as a receive-only neighbor (connected, summary, and static routes are not sent):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# stub receive-only
```

Related Commands

| Command | Description |
|--|---|
| redistribute (EIGRP), on page 655 | Redistributes routes from one routing domain into EIGRP. |
| summary-address (EIGRP), on page 693 | Configures a summary aggregate address for the specified EIGRP interface. |
| auto-summary (EIGRP), on page 609 | Allows automatic summarization of subnet routes into network-level routes for an EIGRP process. |

summary-address (EIGRP)

To configure a summary aggregate address for the specified Enhanced Interior Gateway Routing Protocol (EIGRP) interface, use the **summary-address** command in interface configuration mode. To disable a configuration, use the **no** form of this command.

summary-address *ip-address* {/length| mask} [*admin-distance*]

no summary-address *ip-address* {/length| mask}

Syntax Description

| | |
|-----------------------|--|
| <i>ip-address</i> | The IP address argument specifies the summary IP address to apply to an interface in four-part, dotted-decimal notation. |
| / <i>length</i> | Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If <i>length</i> is used, the slash is required. |
| <i>mask</i> | IP address mask. |
| <i>admin-distance</i> | (Optional) Administrative distance. A value from 1 to 255. |

Command Default

An administrative distance of 5 is applied to EIGRP summary routes.
No summary addresses are predefined.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **summary-address** command is used to configure interface-level address summarization. EIGRP summary routes are given an administrative distance of 5. The administrative distance is used to advertise a summary without installing it in the routing table.

By default, EIGRP summarizes subnet routes to the network level. The **no auto-summary** command can be entered to configure subnet level summarization.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | eigrp | read, write |

Examples

The following example shows how to configure an administrative distance of 95 on an EIGRP interface for the 192.168.0.0/16 summary address:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-eigrp-af-if)# summary-address 192.168.0.0/16 95
```

Related Commands

| Command | Description |
|---|---|
| auto-summary (EIGRP), on page 609 | Allows automatic summarization of subnet routes into network-level routes for an EIGRP process. |

timers active-time

To adjust the Enhanced Interior Gateway Routing Protocol (EIGRP) routing wait time, use the **timers active-time** command in the appropriate configuration mode. To disable this function, use the **timers active-time no** form of the command.

timers active-time [*time-limit*] **disabled**

no timers active-time

Syntax Description

| | |
|-------------------|--|
| <i>time-limit</i> | Active time limit (in minutes). Range is from 1 to 4294967295 minutes. |
| disabled | Disables the timers and permits the routing wait time to remain active indefinitely. |

Command Default

Disabled

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers active-time** command to control the time the router waits (after query is sent) before declaring the route to be in the stuck in active (SIA) state.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to configure an indefinite routing wait time on the specified EIGRP route:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4
```

```
RP/0/RSP0/CPU0:router(config-eigrp-af)# timers active-time disabled
```

Related Commands

| Command | Description |
|--|---|
| show eigrp topology, on page 677 | Displays entries in the EIGRP topology table. |

timers nsf route-hold (EIGRP)

To set the timer that determines how long an NSF-aware Enhanced Interior Gateway Routing Protocol (EIGRP) router holds routes for an inactive peer, use the **timers nsf route-hold** command in the appropriate configuration mode. To return the route hold timer to the default value, use the **no** form of this command.

timers nsf route-hold *seconds*

no timers nsf route-hold

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Time, in seconds, that EIGRP holds routes for an inactive peer. Range is from 20 to 6000 seconds. |
|----------------|---|

Command Default

EIGRP NSF awareness is enabled.

seconds :480

Command Modes

IPv4 address family configuration

IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.1 | The default route hold time was changed from 240 seconds to 480 seconds. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers nsf route-hold** command to set the maximum period of time that the NSF-aware router holds known routes for an NSF-capable neighbor during a switchover operation or a well-known failure condition. The route hold timer is configurable so that you can tune network performance and avoid undesired effects, such as packet loss on routes if the switchover operation takes too much time. When this timer expires, the NSF-aware router scans the topology table and discards any stale routes, allowing EIGRP peers to find alternate routes instead of waiting during a long switchover operation.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set the route hold timer value for an NSF-aware router to 2 minutes (120 seconds):

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af) # timers nsf route-hold 120
```

variance

To control load balancing in an Enhanced Interior Gateway Routing Protocol (EIGRP)-based internetwork, use the **variance** command in the appropriate configuration mode. To reset the variance to the default value, use the **no** form of this command.

variance *multiplier*

no variance

Syntax Description

| | |
|-------------------|---|
| <i>multiplier</i> | Metric value used for load balancing. Range is from 1 to 128. |
|-------------------|---|

Command Default

multiplier : 1 (equal-cost load balancing)

Command Modes

IPv4 address family configuration
IPv4 VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **variance** command to set a variance on the EIGRP router so that the router can determine the feasibility of a potential route. A route is feasible if the next router in the path is closer to the destination than the current router and the metric for the entire path is within the variance. Only paths that are feasible can be used for load balancing and included in the routing table.

If the following two conditions are met, the route is considered feasible and can be added to the routing table:

- 1 The local best metric must be greater than the metric learned from the next router.
- 2 The multiplier times the local best metric for the destination must be greater than or equal to the metric through the next router.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to set a variance of 4:

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1  
RP/0/RSP0/CPU0:router(config-eigrp) address-family ipv4  
RP/0/RSP0/CPU0:router(config-eigrp-af) # variance 4
```

vrf (EIGRP)

To define a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove a VRF instance, use the **no** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

| | |
|-----------------|--------------------------------------|
| <i>vrf-name</i> | VPN routing and forwarding instance. |
|-----------------|--------------------------------------|

Command Default

No VRFs are defined.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

From VRF configuration mode, you must enter address family configuration mode and then issue commands, such as the **auto-summary** command.

A single EIGRP routing process can support multiple VRFs. The number of VRFs that can be configured is limited by available system resources on the router, which is determined by the number of VRFs, running processes, and available memory. However, only a single VRF can be supported by each VPN. Redistribution between different VRFs is not supported.

MPLS VPN support between PE and customer edge (CE) routers is configured only on PE routers that provide VPN services over the service provider backbone. The customer site does not require any changes to equipment or configurations to support the EIGRP VPN. Typically, a metric must be configured for routes to be advertised to the CE router. The metric can be configured under the route-policy for the **redistribute** protocol command or configured with the **default-metric** command.

You must remove IPv4/IPv6 addresses from an interface prior to assigning, removing, or changing a VRF on an IP interface. If this is not done in advance, any attempt to change the VRF on an IP interface is rejected.

Task ID

| Task ID | Operations |
|---------|-------------|
| eigrp | read, write |

Examples

The following example shows how to enter IPv4 VRF address family configuration mode and identify EIGRP commands that can be issued from that mode.

```
RP/0/RSP0/CPU0:router(config)# router eigrp 1
RP/0/RSP0/CPU0:router(config-eigrp)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-eigrp-vrf)# address-family ipv4
RP/0/RSP0/CPU0:router(config-eigrp-vrf-af)# ?

    auto-summary          Auto summarisation
    autonomous-system     Set the autonomous system of VRF
    commit                 Commit the configuration changes to running
    default-information    Handling of default route
    default-metric         Set metric of redistributed routes
    describe               Describe a command without taking real actions
    distance               Set distance for EIGRP routes
    do                     Run an exec command
    exit                   Exit from this submode
    interface              EIGRP interface configuration submode
    log-neighbor-changes   Enable/Disable EIGRP neighbor logging
    log-neighbor-warnings  Enable/Disable EIGRP neighbor warnings
    maximum-paths          Maximum paths
    maximum-prefix         Maximum number of IP prefixes acceptable in aggregate
    metric                 Modify EIGRP routing metrics and parameters
    neighbor               Neighbor prefix limits configuration
    no                     Negate a command or set its defaults
    redistribute            Redistribute another protocol
    route-policy           Configure inbound/outbound policies
    router-id              Set router ID
    show                   Show contents of configuration
    stub                   EIGRP stub
    timers                 Configure EIGRP timers
    variance               Control load balancing variance
```

Related Commands

| Command | Description |
|---|---|
| default-metric (EIGRP), on page 621 | Sets metrics for an EIGRP. |
| redistribute (EIGRP), on page 655 | Injects routes from one routing domain EIGRP. |



IS-IS Commands

This module describes the commands used to configure and monitor the Intermediate System-to-Intermediate System (IS-IS) protocol on Cisco ASR 9000 Series Aggregation Services Routers .

For detailed information about IS-IS concepts, configuration tasks, and examples, see the *Implementing IS-IS on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family \(IS-IS\), page 707](#)
- [address-family multicast topology \(IS-IS\), page 709](#)
- [adjacency-check disable, page 711](#)
- [circuit-type, page 713](#)
- [clear isis process, page 715](#)
- [clear isis route, page 717](#)
- [clear isis statistics, page 719](#)
- [csnp-interval, page 721](#)
- [default-information originate \(IS-IS\), page 723](#)
- [disable \(IS-IS\), page 725](#)
- [distance \(IS-IS\), page 726](#)
- [fast-reroute per-link \(IS-IS\), page 728](#)
- [fast-reroute per-prefix \(IS-IS\), page 730](#)
- [fast-reroute per-link priority-limit \(IS-IS\), page 732](#)
- [fast-reroute per-prefix load-sharing disable \(IS-IS\), page 734](#)
- [fast-reroute per-prefix tiebreaker \(IS-IS\), page 735](#)
- [hello-interval \(IS-IS\), page 737](#)
- [hello-multiplier, page 739](#)
- [hello-padding, page 741](#)
- [hello-password, page 743](#)

- [hello-password keychain](#), page 745
- [hello-password accept](#), page 747
- [hostname dynamic disable](#), page 749
- [ignore-lsp-errors](#), page 751
- [interface \(IS-IS\)](#), page 753
- [ispf](#), page 755
- [is-type](#), page 756
- [log adjacency changes \(IS-IS\)](#), page 758
- [log pdu drops](#), page 760
- [lsp fast-flood threshold](#), page 761
- [lsp-gen-interval](#), page 763
- [lsp-interval](#), page 765
- [lsp-mtu](#), page 767
- [lsp-password](#), page 769
- [lsp-password accept](#), page 771
- [lsp-refresh-interval](#), page 773
- [maximum-paths \(IS-IS\)](#), page 775
- [maximum-redistributed-prefixes \(IS-IS\)](#), page 776
- [max-lsp-lifetime](#), page 778
- [mesh-group \(IS-IS\)](#), page 780
- [metric \(IS-IS\)](#), page 782
- [metric-style narrow](#), page 784
- [metric-style transition](#), page 786
- [metric-style wide](#), page 788
- [min-lsp-arrivaltime](#), page 790
- [mpls ldp auto-config](#), page 792
- [mpls ldp sync \(IS-IS\)](#), page 794
- [mpls traffic-eng \(IS-IS\)](#), page 796
- [mpls traffic-eng multicast-intact \(IS-IS\)](#), page 798
- [mpls traffic-eng path-selection ignore overload](#), page 800
- [mpls traffic-eng router-id \(IS-IS\)](#), page 802
- [net](#), page 804
- [nsf \(IS-IS\)](#), page 806

- [nsf interface-expires](#), page 808
- [nsf interface-timer](#), page 810
- [nsf lifetime \(IS-IS\)](#), page 812
- [passive \(IS-IS\)](#), page 813
- [point-to-point](#), page 814
- [priority \(IS-IS\)](#), page 815
- [propagate level](#), page 817
- [redistribute \(IS-IS\)](#), page 819
- [retransmit-interval \(IS-IS\)](#), page 823
- [retransmit-throttle-interval](#), page 825
- [router isis](#), page 827
- [set-attached-bit](#), page 829
- [set-overload-bit](#), page 831
- [show isis](#), page 833
- [show isis adjacency](#), page 836
- [show isis adjacency-log](#), page 839
- [show isis checkpoint adjacency](#), page 841
- [show isis checkpoint interface](#), page 843
- [show isis checkpoint lsp](#), page 845
- [show isis database](#), page 847
- [show isis database-log](#), page 856
- [show isis fast-reroute](#), page 859
- [show isis hostname](#), page 861
- [show isis interface](#), page 863
- [show isis lsp-log](#), page 868
- [show isis mesh-group](#), page 871
- [show isis mpls traffic-eng adjacency-log](#), page 873
- [show isis mpls traffic-eng advertisements](#), page 875
- [show isis mpls traffic-eng tunnel](#), page 878
- [show isis neighbors](#), page 880
- [show isis protocol](#), page 884
- [show isis route](#), page 887
- [show isis spf-log](#), page 890

- [show isis statistics](#), page 898
- [show isis topology](#), page 902
- [show protocols \(IS-IS\)](#), page 905
- [shutdown \(IS-IS\)](#), page 909
- [single-topology](#), page 910
- [snmp-server traps isis](#), page 912
- [spf-interval](#), page 913
- [spf prefix-priority \(IS-IS\)](#), page 915
- [summary-prefix \(IS-IS\)](#), page 917
- [suppressed](#), page 919
- [tag \(IS-IS\)](#), page 921
- [topology-id](#), page 923
- [trace \(IS-IS\)](#), page 925

address-family (IS-IS)

To enter address family configuration mode for configuring Intermediate System-to-Intermediate System (IS-IS) routing that use standard IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes, use the **address-family** command in router configuration or interface configuration mode. To disable support for an address family, use the **no** form of this command.

address-family {**ipv4**|**ipv6**} {**unicast**|**multicast**}

no address-family {**ipv4**|**ipv6**} {**unicast**|**multicast**}

Syntax Description

| | |
|------------------|---------------------------------------|
| ipv4 | Specifies IPv4 address prefixes. |
| ipv6 | Specifies IPv6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |

Command Default

An address family is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address family** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can configure routing that uses standard IPv4 or IPv6 address prefixes. An address family must be specified in interface configuration mode. In interface address family configuration mode, you can alter interface parameters for IPv4 or IPv6.

You must specify an address family in order to configure parameters that pertain to a single address family.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| isis | read, write |

Examples

The following example shows how to configure the IS-IS router process with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# interface gigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-if-af)#
```

address-family multicast topology (IS-IS)

To enable a multicast topology when configuring Intermediate System-to-Intermediate System (IS-IS) routing (or to place a given topology within the IS-IS interface), use the **address-family multicast topology** command with either IPv4 or IPv6 address prefix in the appropriate configuration mode. To disable a multicast topology in IS-IS, use the **no** form of this command.

address-family {**ipv4**|**ipv6**} **multicast topology** *topo-name* [**maximum prefix** *prefix-limit*]

no address-family {**ipv4**|**ipv6**} **multicast topology** *topo-name*

Syntax Description

| | |
|----------------------------------|---|
| ipv4 | Specifies IPv4 address prefixes. |
| ipv6 | Specifies IPv6 address prefixes. |
| multicast | Specifies multicast address prefixes. |
| topology <i>topo-name</i> | Specifies the name of the topology. |
| maximum prefix | Specifies maximum number of prefixes that a routing table can have. |
| <i>prefix-limit</i> | Maximum number of prefixes. Range is from 32 to 2,000,000. |

Command Default

An address family for multicast topology is not specified. The default subaddress family (SAFI) is unicast.

Command Modes

Router configuration
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address family multicast topology** command to place the router or interface in address family configuration mode. In router address family configuration mode, you can associate an IS-IS topology ID with the topology you have created to add connected and local routes to a specific routing table.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the IS-IS router topology with an IPv4 multicast address prefix:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 multicast topology green
RP/0/RSP0/CPU0:router(config-isis-af)#
```

or

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface gigabitethernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 multicast topology green
RP/0/RSP0/CPU0:router(config-isis-if-af)#
```

Related Commands

| Command | Description |
|---|---|
| topology-id , on page 923 | Associates a topology ID with a named IS-IS topology to differentiate topologies in the domain. |

adjacency-check disable

To suppress Intermediate System-to-Intermediate System (IS-IS) IP Version 4 (IPv4) or IP Version 6 (IPv6) protocol-support consistency checks that are performed prior to forming adjacencies on hello packets, use the **adjacency-check disable** command in address family configuration mode. To remove this function, use the **no** form of this command.

adjacency-check disable

no adjacency-check disable

Command Default Adjacency check is enabled

Command Modes Address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |
| | Release 3.9.0 | Support was added for ipv6. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS performs consistency checks on hello packets and forms an adjacency only with a neighboring router that supports the same set of protocols. A router running IS-IS for both IPv4 and IPv6 does not form an adjacency with a router running IS-IS for IPv4 only.

Use the **adjacency-check disable** command to suppress the consistency checks for IPv6 IS-IS and allow an IPv4 IS-IS router to form an adjacency with a router running IPv4 IS-IS and IPv6. IS-IS never forms an adjacency between a router running IPv4 IS-IS only and a router running IPv6 only.

In addition, the **adjacency-check disable** command suppresses the IPv4 or IPv6 subnet consistency check and allows IS-IS to form an adjacency with other routers regardless of whether they have an IPv4 or IPv6 subnet in common.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The command in the following example disables the adjacency checks:

The following example shows how the network administrator introduces IPv6 into an existing IPv4 IS-IS network and ensures that the checking of hello packet checks from adjacent neighbors is disabled until all neighbor routers are configured to use IPv6:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 |ipv4  
RP/0/RSP0/CPU0:router(config-isis-af)# adjacency-check disable
```

circuit-type

To configure the type of adjacency used for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **circuit-type** command in interface configuration mode. To reset the circuit type to Level 1 and Level 2, use the **no** form of this command.

```
circuit-type {level-1| level-1-2| level-2-only}
no circuit-type
```

Syntax Description

| | |
|---------------------|--|
| level-1 | Establishes only Level 1 adjacencies over an interface. |
| level-1-2 | Establishes both Level 1 and Level 2 adjacencies, if possible. |
| level-2-only | Establishes only Level 2 adjacencies over an interface. |

Command Default

Default adjacency types are Level 1 and Level 2 adjacencies.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Adjacencies may not be established even if allowed by the **circuit-type** command. The proper way to establish adjacencies is to configure a router as a Level 1, Level 1 and Level 2, or Level 2-only system using the [is-type, on page 756](#) command. Only on networking devices that are between areas (Level 1 and Level 2 networking devices) should you configure some interfaces to be Level 2-only to prevent wasting bandwidth by sending out unused Level 1 hello packets. Remember that on point-to-point interfaces, the Level 1 and Level 2 hello packets are in the same packet.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a Level 1 adjacency with its neighbor on GigabitEthernet interface 0/2/0/0 and Level 2 adjacencies with all Level 2-capable routers on GigabitEthernet interface 0/5/0/2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-1-2
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-1
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/5/0/2
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-2-only
```

In this example, only Level 2 adjacencies are established because the **is-type** command is configured:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-2-only
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# circuit-type level-1-2
```

Related Commands

| Command | Description |
|---------------------------------------|--|
| is-type , on page 756 | Configures the routing level for an instance of the IS-IS routing process. |
| net , on page 804 | Configures an IS-IS NET for the routing process. |

clear isis process

To clear the link-state packet (LSP) database and adjacency database sessions for an Intermediate System-to-Intermediate System (IS-IS) instance or all IS-IS instances, use the **clear isis process** command in EXEC configuration mode.

clear isis [**instance** *instance-id*] **process**

Syntax Description

instance *instance-id* (Optional) Specifies IS-IS sessions for the specified IS-IS instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear isis process** command without any keyword to clear all the IS-IS instances. Add the **instance** *instance-id* keyword and argument to clear the specified IS-IS instance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows the IS-IS LSP database and adjacency sessions being cleared for instance 1:

```
RP/0/RSP0/CPU0:router# clear isis instance 1 process
```

Related Commands

| Command | Description |
|--|---|
| show isis database, on page 847 | Displays the IS-IS link-state database. |
| show isis neighbors, on page 880 | Displays information about IS-IS neighbors. |

clear isis route

To clear the Intermediate System-to-Intermediate System (IS-IS) routes in a topology, use the **clear isis route** command in EXEC configuration mode.

```
clear isis [instance instance-id] {afi-all| ipv4| ipv6} {unicast| multicast| safi-all} [topology topo-name]
route
```

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Specifies IS-IS sessions for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| afi-all | Specifies IP Version 4 (IPv4) and IP Version 6 (IPv6) address prefixes. |
| ipv4 | Specifies IPv4 address prefixes. |
| ipv6 | Specifies IPv6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| safi-all | Specifies all secondary address prefixes. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear isis route** command to clear the routes from the specified topology or all routes in all topologies if no topology is specified.

Task ID

| Task ID | Operations |
|----------------|-------------|
| isis | execute |
| rib | read, write |
| basic-services | read, write |

Examples

The following example shows how to clear the routes with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router# clear isis ipv4 unicast route
```

Related Commands

| Command | Description |
|--|---|
| show isis database, on page 847 | Displays the IS-IS link-state database. |
| show isis neighbors, on page 880 | Displays information about IS-IS neighbors. |

clear isis statistics

To clear the Intermediate System-to-Intermediate System (IS-IS) statistics, use the **clear isis statistics** command in EXEC configuration mode.

clear isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Clears IS-IS sessions for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear isis statistics** command to clear the information displayed by the **show isis statistics** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | execute |
| rib | read, write |

| Task ID | Operations |
|----------------|-------------|
| basic-services | read, write |

Examples

The following example shows the IS-IS statistics for a specified interface being cleared:

```
RP/0/RSP0/CPU0:router# clear isis instance 23 statistics
```

Related Commands

| Command | Description |
|---|--------------------------------|
| show isis statistics, on page 898 | Displays the IS-IS statistics. |

csnp-interval

To configure the interval at which periodic complete sequence number PDU (CSNP) packets are sent on broadcast interfaces, use the **csnp-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

csnp-interval *seconds* [**level** {**1** | **2**}]

no csnp-interval *seconds* [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|---|
| <i>seconds</i> | Interval (in seconds) of time between transmission of CSNPs on multiaccess networks. This interval applies only for the designated router. Range is 0 to 65535 seconds. |
| level { 1 2 } | (Optional) Specifies the interval of time between transmission of CSNPs for Level 1 or Level 2 independently. |

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **csnp-interval** command applies only to the designated router (DR) for a specified interface. Only DRs send CSNP packets to maintain database synchronization. The CSNP interval can be configured independently for Level 1 and Level 2.

Use of the **csnp-interval** command on point-to-point subinterfaces makes sense only in combination with the IS-IS mesh-group feature.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | execute |
| rib | read, write |

| Task ID | Operations |
|----------------|-------------|
| basic-services | read, write |

Examples

The following example shows how to set the CSNP interval for Level 1 to 30 seconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/0/2/0  
RP/0/RSP0/CPU0:router(config-isis-if)# cstp-interval 30 level 1
```

default-information originate (IS-IS)

To generate a default route into an Intermediate System-to-Intermediate System (IS-IS) routing domain, use the **default-information originate** command in address family configuration mode. To remove the **default-information originate** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

default-information originate [**route-policy** *route-policy-name*]

no default-information originate [**external**] **route-policy** *route-policy-name*]

Syntax Description

| | |
|--------------------------|--|
| route-policy | (Optional) Defines the conditions for the default route. |
| <i>route-policy-name</i> | (Optional) Name for the route policy. |

Command Default

A default route is not generated into an IS-IS routing domain.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If a router configured with the **default-information originate** command has a route to 0.0.0.0 in the routing table, IS-IS originates an advertisement for 0.0.0.0 in its link-state packets (LSPs).

Without a route policy, the default is advertised only in Level 2 LSPs. For Level 1 routing, there is another process to find the default route, which is to look for the closest Level 1 and Level 2 router. The closest Level 1 and Level 2 router can be found by looking at the attached-bit (ATT) in Level 1 LSPs.

A route policy can be used for two purposes:

- To make the router generate the default route in its Level 1 LSPs.
- To advertise 0.0.0.0/0 conditionally.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to generate a default external route into an IS-IS domain:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# default-information originate
```

Related Commands

| Command | Description |
|---|---|
| redistribute (IS-IS), on page 819 | Redistributes routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS). |
| show isis database, on page 847 | Displays the IS-IS link-state database. |

disable (IS-IS)

To disable the Intermediate System-to-Intermediate System (IS-IS) topology on a specified interface, use the **disable** command in interface address family configuration mode. To remove this function, use the **no** form of this command.

disable

no disable

Command Default IS-IS protocol is enabled.

Command Modes Interface address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to disable the IS-IS protocol for IPv4 unicast on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# disable
```

distance (IS-IS)

To define the administrative distance assigned to routes discovered by the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **distance** command in address family configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*prefix mask* | *prefix/length* | [*prefix-list-name*]]

no distance [*weight*] [*prefix mask* | *prefix/length* | [*prefix-list-name*]]

Syntax Description

| | |
|-------------------------|--|
| <i>weight</i> | Administrative distance to be assigned to IS-IS routes. Range is 1 to 255. |
| <i>prefix</i> | (Optional) The <i>prefix</i> argument specifies the IP address in four-part, dotted-decimal notation. |
| <i>mask</i> | (Optional) IP address mask. |
| <i>/length</i> | (Optional) The length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. Range is 0 to 32 for IPv4 addresses and 0 to 128 for IPv6 addresses. |
| <i>prefix-list-name</i> | (Optional) List of routes to which administrative distance applies. |

Command Default

weight : 115

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An administrative distance is an integer from 1 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

Use the **distance** command to configure the administrative distances applied to IS-IS routes when they are inserted into the Routing Information Base (RIB), and influence the likelihood of these routes being preferred over routes to the same destination addresses discovered by other protocols.

The *address/prefix-length* argument defines to which source router the distance applies. In other words, each IS-IS route is advertised by another router, and that router advertises an address that identifies it. This source address is displayed in the output of the **show isis route detail** command.

The **distance** command applies to the routes advertised by routers whose address matches the specified prefix. The *prefix-list-name* argument can then be used to refine this further so that the **distance** command affects only specific routes.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

In the following example, a distance of 10 is assigned to all routes to 2.0.0.0/8 and 3.0.0.0/8 (or more specific prefixes) that are advertised by routers whose ID is contained in 1.0.0.0/8. A distance of 80 is assigned to all other routes.

```
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list target_routes
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# permit 2.0.0.0/8
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# permit 3.0.0.0/8
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# deny 0.0.0.0/0
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# exit
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# distance 10 1.0.0.0/8 target_routes
RP/0/RSP0/CPU0:router(config-isis-af)# distance 80
```

Related Commands

| Command | Description |
|---|--|
| router isis , on page 827 | Configures the IS-IS routing protocol and specifies an IS-IS instance. |
| show isis protocol , on page 884 | Displays summary information about the IS-IS instance. |
| show isis route , on page 887 detail | Displays link-state packet (LSP) details. |

fast-reroute per-link (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link** command in interface address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link [**exclude interface** *type interface-path-id*] **level** {1|2} | **lfa-candidate interface** *type interface-path-id*]

no fast-reroute per-link

Syntax Description

| | |
|--------------------------|--|
| exclude | Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information |
| level {1 2} | Configures FRR LFA computation for one level only. |
| lfa-candidate | Specifies FRR LFA computation candidate information |
| interface | Specifies an interface that needs to be either excluded from FRR LFA computation (when used with exclude keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the lfa-candidate keyword). |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

IP fast-reroute LFA per-link computation is disabled.

Command Modes

Interface address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to configure per-link fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# interface POS0/3/0/0  
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-if-af)# fast-reroute per-link level 1
```

Related Commands

| Command | Description |
|--|---|
| fast-reroute per-prefix (IS-IS), on page 730 | Enables IP fast reroute (IPFRR) loop-free alternate (LFA) prefix dependent computation. |

fast-reroute per-prefix (IS-IS)

To enable IP fast reroute (IPFRR) loop-free alternate (LFA) prefix dependent computation, use the **fast-reroute per-prefix** command in interface address family configuration mode. LFA is supported only on Enhanced Ethernet line card. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix [**exclude interface** *type interface-path-id* | **level** {1| 2}| **lfa-candidate interface** *type interface-path-id* | **remote-lfa** {**maximum-metric** *metric-value* | **tunnel mpls-ldp**} [**level** {1| 1}]]

no fast-reroute per-prefix

Syntax Description

| | |
|--------------------------|--|
| exclude | Specifies fast-reroute (FRR) loop-free alternate (LFA) computation exclusion information |
| level {1 2} | Configures FRR LFA computation for one level only. |
| lfa-candidate | Specifies FRR LFA computation candidate information |
| interface | Specifies an interface that needs to be either excluded from FRR LFA computation (when used with exclude keyword) or to be included to LFA candidate list in FRR LFA computation (when used with the lfa-candidate keyword). |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

IP fast-reroute LFA per-prefix computation is disabled.

Command Modes

Interface address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to configure per-prefix fast-reroute LFA computation for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface POS0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# fast-reroute per-prefix level 1
```

Related Commands

| Command | Description |
|--|--|
| fast-reroute per-link (IS-IS), on page 728 | Enables IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation. |

fast-reroute per-link priority-limit (IS-IS)

To enable the IP fast reroute (IPFRR) loop-free alternate (LFA) prefix independent per-link computation, use the **fast-reroute per-link priority-limit** command in address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link priority-limit {critical| high| medium} level {1| 2}

no fast-reroute per-link priority-limit

Syntax Description

| | |
|--------------------|---|
| critical | Enables LFA computation for critical priority prefixes only. |
| high | Enables LFA computation for for criticaland high priority prefixes. |
| medium | Enables LFA computation for for critical, high, and medium priority prefixes. |
| level {1 2} | Sets priority-limit for routing Level 1 or Level 2 independently. |

Command Default

Fast-reroute per link priority limit LFA computation is disabled.

Command Modes

IPv4 unicast address family configuration
 IPv4 multicast address family configuration
 IPv6 unicast address family configuration
 IPv6 multicast address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to configure fast-reroute prefix independent per-link computation for critical priority prefixes for level 1 only:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-link priority-limit critical level
1
```

fast-reroute per-prefix load-sharing disable (IS-IS)

To disable load sharing prefixes across multiple backups, use the **fast-reroute per-prefix load-sharing disable** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix load-sharingdisable

no fast-reroute per-prefix load-sharingdisable

Syntax Description

| | |
|--------------------|---|
| level {1 2} | Disables load-sharing for Level 1 or Level 2 independently. |
|--------------------|---|

Command Default

Load sharing is enabled.

Command Modes

IPv4 unicast address family configuration

IPv4 multicast address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to disable load-sharing prefixes across multiple backups for level 1 routes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix load-sharing disable level 1
```

fast-reroute per-prefix tiebreaker (IS-IS)

To configure tie-breaker for multiple backups, use the **fast-reroute per-prefix tiebreaker** command in IPv4 address family configuration mode. To disable tie-breaker configuration, use the **no** form of this command.

fast-reroute per-prefix tiebreaker [**downstream** | **lc-disjoint** | **lowest-backup-metric** | **node-protecting** | **primary-path** | **secondary-path**] **index** *index_number* **level** {**1** | **2**}

no fast-reroute per-prefix tiebreaker

Syntax Description

| | |
|--------------------------------------|---|
| downstream | Configures to prefer backup path via downstream node, in case of tie-breaker. |
| lc-disjoint | Configures to prefer Prefer line card disjoint backup path. |
| lowest-backup-metric | Configures to prefer backup path with lowest total metric. |
| node-protecting | Configures to prefer node protecting backup path. |
| primary-path | Configures to prefer backup path from ECMP set. |
| secondary-path | Configures to prefer non-ECMP backup path. |
| index | Sets preference order among tie-breakers. |
| <i>index_number</i> | Value for the index. Range is 1-255. |
| level { 1 2 } | Configures tiebreaker for Level 1 or Level 2 independently. |

Command Default

Tie-breaker for multiple backups is not configured.

Command Modes

IPv4 unicast address family configuration
IPv4 multicast address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples

This example shows how to configure preference of backup path via downstream node in case of a tie-breaker for selection of backup path from multiple backup paths:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp_lfa
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4
RP/0/RSP0/CPU0:router(config-isis-af)#fast-reroute per-prefix tiebreaker downstream index
255
```

hello-interval (IS-IS)

To specify the length of time between consecutive hello packets sent by the Intermediate System-to-Intermediate System (IS-IS) protocol software, use the **hello-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-interval *seconds* [**level** {**1** | **2**}]

no hello-interval [*seconds*] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|--|
| <i>seconds</i> | Integer value (in seconds) for the length of time between consecutive hello packets. By default, a value three times the hello interval <i>seconds</i> is advertised as the <i>hold time</i> in the hello packets sent. (That multiplier of three can be changed by using the hello-multiplier command.) With smaller hello intervals, topological changes are detected more quickly, but there is more routing traffic. Range is 1 to 65535 seconds. |
| level { 1 2 } | (Optional) Specifies the hello interval for Level 1 and Level 2 independently. For broadcast interfaces only. |

Command Default

seconds : 10 seconds

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval can be configured independently for Level 1 and Level 2, except on serial point-to-point interfaces. (Because only a single type of hello packet is sent on serial links, it is independent of Level 1 or Level 2.) Configuring Level 1 and Level 2 independently is used on LAN interfaces.



Note

A shorter hello interval gives quicker convergence, but increases bandwidth and CPU usage. It might also add to instability in the network.

A slower hello interval saves bandwidth and CPU. Especially when used in combination with a higher hello multiplier, this strategy may increase overall network stability.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure TenGigE interface 0/6/0/0 to advertise hello packets every 5 seconds for Level 1 topology routes. This situation causes more traffic than configuring a longer interval, but topological changes are detected more quickly.

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# hello-interval 5 level 1
```

Related Commands

| Command | Description |
|---|--|
| hello-multiplier, on page 739 | Specifies the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down. |

hello-multiplier

To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the **hello-multiplier** command in interface configuration mode. To restore the default value, use the **no** form of this command.

hello-multiplier *multiplier* [**level** {**1** | **2**}]

no hello-multiplier [*multiplier*] [**level** {**1** | **2**}]

Syntax Description

multiplier Advertised hold time in IS-IS hello packets is set to the hello multiplier times the hello interval. Range is 3 to 1000. Neighbors declare an adjacency to this down router after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on an individual interface basis, and can be different between different networking devices in one area.

Using a smaller hello multiplier gives faster convergence, but can result in more routing instability. Increase the hello multiplier to a larger value to help network stability when needed. Never configure a hello multiplier to a value lower than the default value of 3.

level { **1** | **2** } (Optional) Specifies the hello multiplier independently for Level 1 or Level 2 adjacencies.

Command Default

multiplier : 3

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The “holding time” carried in an IS-IS hello packet determines how long a neighbor waits for another hello packet before declaring the neighbor to be down. This time determines how quickly a failed link or neighbor is detected so that routes can be recalculated.

Use the **hello-multiplier** command in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval ([hello-interval \(IS-IS\)](#), on page 737 command) correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.

On point-to-point links, there is only one hello for both Level 1 and Level 2. Separate Level 1 and Level 2 hello packets are also sent over nonbroadcast multiaccess (NBMA) networks in multipoint mode, such as X.25, Frame Relay, and ATM.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how the network administrator wants to increase network stability by making sure an adjacency goes down only when many (ten) hello packets are missed. The total time to detect link failure is 60 seconds. This strategy ensures that the network remains stable, even when the link is fully congested.

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# hello-interval 6
RP/0/RSP0/CPU0:router(config-isis-if)# hello-multiplier 10
```

Related Commands

| Command | Description |
|---|---|
| hello-interval (IS-IS), on page 737 | Specifies the length of time between hello packets that the software sends. |

hello-padding

To configure padding on Intermediate System-to-Intermediate System (IS-IS) hello protocol data units (IIH PDUs) for all IS-IS interfaces on the router, use the **hello-padding** command in interface configuration mode. To suppress padding, use the **no** form of this command.

hello-padding {disable| sometimes} [level {1| 2}]

no hello-padding {disable| sometimes} [level {1| 2}]

Syntax Description

| | |
|------------------------|--|
| disable | Suppresses hello padding. |
| sometimes | Enables hello padding during adjacency formation only. |
| level { 1 2 } | (Optional) Specifies hello padding for Level 1 or Level 2 independently. |

Command Default

Hello padding is enabled.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might want to suppress hello padding to conserve network resources. The lower the circuit speed, the higher the percentage of padding overhead. Before suppressing the hello padding, you should know your physical and data link layer configurations and have control over them, and also know your router configuration at the network layer.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to suppress IS-IS hello padding over local area network (LAN) circuits for interface GigabitEthernet 0/2/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1  
RP/0/RSP0/CPU0:router(config-isis-if)# hello-padding disable
```

Related Commands

| Command | Description |
|--|---|
| show isis interface, on page 863 | Displays information about the IS-IS interface. |

hello-password

To configure the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password** command in interface configuration mode. To disable authentication, use the **no** form of this command.

```
hello-password [hmac-md5] text [clear| encrypted] password [level {1| 2}] [send-only]
no hello-password [hmac-md5] text [clear| encrypted] password [level {1| 2}] [send-only]
```

Syntax Description

| | |
|------------------------|---|
| hmac-md5 | (Optional) Specifies that the password use HMAC-MD5 authentication. |
| text | (Optional) Specifies that the password use clear text password authentication. |
| clear | (Optional) Specifies that the password be unencrypted. |
| encrypted | (Optional) Specifies that the password be encrypted using a two-way algorithm. |
| <i>password</i> | Authentication password you assign for an interface. |
| level { 1 2 } | (Optional) Specifies whether the password is for a Level 1 or a Level 2 protocol data unit (PDU). |
| send-only | (Optional) Specifies that the password applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a **text** password is configured, it is exchanged as clear text. Therefore, the **hello-password** command provides limited security.

When an **hmac-md5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

For point-to-point links, IS-IS sends only a single hello for Level 1 and Level 2, making the **level** keyword meaningless on point-to-point links. To modify hello parameters for a point-to-point interface, omit the **level** keyword.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a password with HMAC-MD5 authentication for hello packets running on GigabitEthernet 0/2/0/3 interface:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-isis-if)# hello-password hmac-md5 clear mypassword
```

Related Commands

| Command | Description |
|--|--|
| hello-password keychain, on page 745 | Configures the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface. |
| hello-password accept, on page 747 | Configures an additional authentication password for an IS-IS interface. |

hello-password keychain

To configure the authentication password keychain for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password keychain** command in interface configuration mode. To disable the authentication password keychain, use the **no** form of this command.

hello-password keychain *keychain-name* [**level** {1|2}] [**send-only**]

no hello-password keychain *keychain-name* [**level** {1|2}] [**send-only**]

Syntax Description

| | |
|------------------------|---|
| keychain | Keyword that specifies the keychain to be configured. An authentication password keychain is a sequence of keys that are collectively managed and used for authenticating a peer-to-peer group. |
| <i>keychain-name</i> | Specifies the name of the keychain. |
| level { 1 2 } | (Optional) Specifies whether the keychain is for a Level 1 or a Level 2 protocol data unit (PDU). |
| send-only | (Optional) Specifies that the keychain applies only to protocol data units (PDUs) that are being sent and does not apply to PDUs that are being received. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

password: encrypted text

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Specify a keychain to enable keychain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a password keychain for level 1, send only authentication on a GigabitEthernet interface:

```
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0

RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis-if)# hello-password keychain
mykeychain level 1 send-only
```

Related Commands

| Command | Description |
|---|---|
| hello-password , on page 743 | Configures the authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface. |
| hello-password accept , on page 747 | Configures an additional authentication password for an IS-IS interface. |

hello-password accept

To configure an additional authentication password for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **hello-password accept** command in interface configuration mode. To disable authentication, use the **no** form of this command.

hello-password accept {clear| encrypted} *password* [level {1| 2}]

no hello-password accept {clear| encrypted} *password* [level {1| 2}]

Syntax Description

| | |
|------------------------|---|
| clear | Specifies that the password be unencrypted. |
| encrypted | Specifies that the password be encrypted using a two-way algorithm. |
| <i>password</i> | Authentication password you assign. |
| level { 1 2 } | (Optional) Specifies the password for Level 1 or Level 2 independently. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **hello-password accept** command to add an additional password for an IS-IS interface. An authentication password must be configured using the **hello-password** command before an accept password can be configured for the corresponding level.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a password:

```
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/3

RP/0/RSP0/CPU0:routerRP/0/RSP0/CPU0:router(config-isis)# hello-password accept encrypted
111D1C1603
```

Related Commands

| Command | Description |
|---|---|
| hello-password, on page 743 | Configures an authentication password for an IS-IS interface. |

hostname dynamic disable

To disable Intermediate System-to-Intermediate System (IS-IS) routing protocol dynamic hostname mapping, use the **hostname dynamic** command in router configuration mode. To remove the specified command from the configuration file and restore the system to its default condition, use the **no** form of this command.

hostname dynamic disable

no hostname dynamic disable

Syntax Description

| | |
|----------------|-------------------------------|
| disable | Disables dynamic host naming. |
|----------------|-------------------------------|

Command Default

Router names are dynamically mapped to system IDs.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

In an IS-IS routing domain, each router is represented by a 6-byte hexadecimal system ID. When network administrators maintain and troubleshoot networking devices, they must know the router name and corresponding system ID.

Link-state packets (LSPs) include the dynamic hostname in the type, length, and value (TLV) which carries the mapping information across the entire domain. Every router in the network, upon receiving the TLV from an LSP, tries to install it in a mapping table. The router then uses the mapping table when it wants to convert a system ID to a router name.

To display the entries in the mapping tables, use the **show isis hostname** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to disable dynamic mapping of hostnames to system IDs:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# hostname dynamic disable
```

Related Commands

| Command | Description |
|---|--|
| hostname | Specifies the name of the local router. |
| show isis hostname, on page 861 | Displays the router name-to-system ID mapping table. |

ignore-lsp-errors

To override the default setting of a router to ignore Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs) that are received with internal checksum errors, use the **ignore-lsp-errors disable** command in router configuration mode. To enable ignoring IS-IS LSP errors, use the **no** form of this command.

ignore-lsp-errors disable

no ignore-lsp-errors disable

Syntax Description

| | |
|----------------|--|
| disable | Disables the functionality of the command. |
|----------------|--|

Command Default

The system purges corrupt LSPs that cause the initiator to regenerate LSPs.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The IS-IS protocol definition requires that a received LSP with an incorrect data-link checksum be purged by the receiver, which causes the initiator of the packet to regenerate it. However, if a network has a link that causes data corruption and at the same time is delivering LSPs with correct data-link checksums, a continuous cycle of purging and regenerating large numbers of packets can occur. Because this situation could render the network nonfunctional, use this command to ignore these LSPs rather than purge the packets.

The receiving network devices use link-state packets to maintain their routing tables.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to instruct the router to ignore LSPs that have internal checksum errors:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# ignore-lsp-errors disable
```


interface (IS-IS)

To configure the Intermediate System-to-Intermediate System (IS-IS) protocol on an interface, use the **interface** command in router configuration mode. To disable IS-IS routing for interfaces, use the **no** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No interfaces are specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An address family must be established on the IS-IS interface before the interface is enabled for IS-IS protocol operation.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to enable an IS-IS multitopology configuration for IPv4 on GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# metric-style wide level 1
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 2001::1/64
```

Related Commands

| Command | Description |
|---|--|
| log adjacency changes (IS-IS) , on page 758 | Configures the routing level for an instance of the IS-IS routing process. |
| net , on page 804 | Configures an IS-IS network entity title (NET) for the routing process. |
| router isis , on page 827 | Enables the IS-IS routing protocol. |

ispf

To configure the incremental shortest path first (iSPF) algorithm to calculate network topology, use the **ispf** command in address family configuration mode. To disable this algorithm function, use the **no** form of this command.

ispf [level {1|2}]

no ispf [level {1|2}]

Syntax Description

| | |
|------------------------|--|
| level { 1 2 } | (Optional) Configures the iSPF algorithm for Level 1 or Level 2 independently. |
|------------------------|--|

Command Default

The iSPF algorithm is not configured.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The iSPF algorithm may be used to reduce the processor load when IS-IS needs to recalculate its topology after minor changes.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure iSPF for the IPv4 unicast topology at Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis ispf
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# ispf level 1
```

is-type

To configure the routing level for an Intermediate System-to-Intermediate System (IS-IS) area, use the **is-type** command in router configuration mode. To set the routing level to the default level, use the **no** form of this command.

is-type {level-1| level-1-2| level-2-only}

no is-type [level-1| level-1-2| level-2-only]

Syntax Description

| | |
|---------------------|--|
| level-1 | Specifies that the router perform only Level 1 (intra-area) routing. This router learns only about destinations inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router. |
| level-1-2 | Specifies that the router perform both Level 1 and Level 2 routing. |
| level-2-only | Specifies that the routing process acts as a Level 2 (interarea) router only. This router is part of the backbone, and does not communicate with Level 1-only routers in its own area. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the router is configured with Level 1 routing only, this router learns about destinations only inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.

When the router is configured with Level 2 routing only, this router is part of the backbone, and does not communicate with Level 1 routers in its own area.

The router has one link-state packet database (LSDB) for destinations inside the area (Level 1 routing) and runs a shortest path first (SPF) calculation to discover the area topology. It also has another LSDB with link-state packets (LSPs) of all other backbone (Level 2) routers, and runs another SPF calculation to discover the topology of the backbone and the existence of all other areas.

We highly recommend that you configure the type of an IS-IS routing process to establish the proper level of adjacencies. If there is only one area in the network, there is no need to run both Level 1 and Level 2 routing algorithms.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to specify that the router is part of the backbone and that it does not communicate with Level 1-only routers:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# is-type level-2-only
```

Related Commands

| Command | Description |
|---|---|
| circuit-type , on page 713 | Configures the type of adjacency. |
| show isis neighbors , on page 880 | Displays information about IS-IS neighbors. |

log adjacency changes (IS-IS)

To cause an IS-IS instance to generate a log message when an Intermediate System-to-Intermediate System (IS-IS) adjacency changes state (up or down), use the **log adjacency changes** command in router configuration mode. To restore the default value, use the **no** form of this command.

log adjacency changes

no log adjacency changes

Command Default No IS-IS instance log messages are generated.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log adjacency changes** command to monitor IS-IS adjacency state changes; it may be very useful when you are monitoring large networks. Messages are logged using the system error message facility. Messages can be in either of two forms:

```
%ISIS-4-ADJCHANGE: Adjacency to 0001.0000.0008 (Gi 0/2/1/0) (L2) Up, new adjacency
%ISIS-4-ADJCHANGE: Adjacency to router-gsr8 (Gi /2/1/0) (L1) Down, Holdtime expired
```

Using the **no** form of the command removes the specified command from the configuration file and restores the system to its default condition with respect to the command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to configure the router to log adjacency changes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# log adjacency changes
```

Related Commands

| Command | Description |
|---------|--|
| logging | Logs messages to a syslog server host. |

log pdu drops

To log Intermediate System-to-Intermediate System (IS-IS) protocol data units (PDUs) that are dropped, use the **log pdu drops** command in router configuration mode. To disable this function, use the **no** form of this command.

log pdu drops

no log pdu drops

Command Default PDU logging is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log pdu drops** command to monitor a network when IS-IS PDUs are suspected of being dropped. The reason for the PDU being dropped and current PDU drop statistics are recorded.

The following are examples of PDU logging output:

```
%ISIS-4-ERR_IIH_INPUT_Q_OVERFLOW: IIH input queue overflow: 86 total drops; 19 IIH drops,
44 LSP drops, 23 SNP drops
%ISIS-4-ERR_LSP_INPUT_Q_OVERFLOW: LSP input queue overflow: 17 total drops; 9 IIH drops,
3 LSP drops, 5 SNP drops
```

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to enable PDU logging:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# log pdu drops
```

lsp fast-flood threshold

To configure the link-state packet (LSP) fast-flood threshold, use the **lsp fast-flood threshold** command in interface configuration mode. To restore the default value, use the **no** form of this command.

lsp fast-flood threshold *lsp-number* [**level** {1|2}]

no lsp fast-flood threshold [*lsp-number*] [**level** {1|2}]

Syntax Description

| | |
|------------------------|--|
| <i>lsp-number</i> | Number of LSPs to send back to back. Range is 1 to 4294967295. |
| level { 1 2 } | (Optional) Specifies the LSP threshold for Level 1 or Level 2 independently. |

Command Default

10 LSPs are allowed in a back-to-back window

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **lsp fast-flood threshold** command to accelerate convergence of LSP database. LSPs are sent back-to-back over an interface up to the specified limit. Past the limit, LSPs are sent out in the next batch window as determined by LSP pacing interval.

Duration of back-to-back window = LSP interval * LSP fast-flood threshold limit.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the LSP threshold:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
```

```
RP/0/RSP0/CPU0:router(config-isis-if)# lsp fast-flood threshold 234 level 1
```

Related Commands

| Command | Description |
|--|--|
| lsp-interval , on page 765 | Configures the amount of time between consecutive LSPs sent on an IS-IS interface. |

lsp-gen-interval

To customize IS-IS throttling of link-state packet (LSP) generation, use the **lsp-gen-interval** command in router configuration mode. To restore the default value, use the **no** form of this command.

lsp-gen-interval [**initial-wait** *initial*] [**secondary-wait** *secondary*] [**maximum-wait** *maximum*] [**level** {**1** | **2**}]
no lsp-gen-interval [[**initial-wait** *initial*] [**secondary-wait** *secondary*] [**maximum-wait** *maximum*]] [**level** {**1** | **2**}]

Syntax Description

| | |
|--|--|
| initial-wait <i>initial</i> | Specifies the initial LSP generation delay (in milliseconds). Range is 0 to 120000 milliseconds. |
| secondary-wait <i>secondary</i> | Specifies the hold time between the first and second LSP generation (in milliseconds). Range is 1 to 120000 milliseconds. |
| maximum-wait <i>maximum</i> | Specifies the maximum interval (in milliseconds) between two consecutive occurrences of an LSP being generated. Range is 1 to 120000 milliseconds. |
| level { 1 2 } | (Optional) Specifies the LSP time interval for Level 1 or Level 2 independently. |

Command Default

initial-wait *initial* : 50 milliseconds
secondary-wait *secondary* : 200 milliseconds
maximum-wait *maximum* : 5000 milliseconds

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

During prolonged periods of network instability, repeated recalculation of LSPs can cause increased CPU load on the local router. Further, the flooding of these recalculated LSPs to the other Intermediate Systems in the network causes increased traffic and can result in other routers having to spend more time running route calculations.

Use the **lsp-gen-interval** command to reduce the rate of LSP generation during periods of instability in the network. This command can help to reduce CPU load on the router and to reduce the number of LSP transmissions to its IS-IS neighbors.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to set the maximum interval between two consecutive occurrences of an LSP to 15 milliseconds and the initial LSP generation delta to 5 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-gen-interval maximum-wait 15 initial-wait 5
```

Related Commands

| Command | Description |
|--|--|
| retransmit-interval (IS-IS), on page 823 | Configures the amount of time between retransmission of each IS-IS LSP on a point-to-point link. |

lsp-interval

To configure the amount of time between consecutive link-state packets (LSPs) sent on an Intermediate System-to-Intermediate System (IS-IS) interface, use the **lsp-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

lsp-interval *milliseconds* [**level** {**1** | **2**}]

no lsp-interval [*milliseconds*] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|---|
| <i>milliseconds</i> | Time delay (in milliseconds) between successive LSPs. Range is 1 to 4294967295. |
| level { 1 2 } | (Optional) Configures the LSP time delay for Level 1 or Level 2 independently. |

Command Default

milliseconds : 33 milliseconds

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to cause the system to send LSPs every 100 milliseconds (10 packets per second) on Level 1 and Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# lsp-interval 100
```

Related Commands

| Command | Description |
|---|--|
| retransmit-interval (IS-IS) , on page 823 | Configures the amount of time between retransmission of each IS-IS LSP on a point-to-point link. |

lsp-mtu

To set the maximum transmission unit (MTU) size of Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs), use the **lsp-mtu** command in router configuration mode. To restore the default, use the **no** form of this command.

lsp-mtu *bytes* [**level** {**1** | **2**}]

no lsp-mtu [*bytes*] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|--|
| <i>bytes</i> | Maximum packet size in bytes. The number of bytes must be less than or equal to the smallest MTU of any link in the network. Range is 128 to 4352 bytes. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Under normal conditions, the default MTU size should be sufficient. However, if the MTU size of a link is less than 1500 bytes, the LSP MTU size must be lowered accordingly on each router in the network. If this action is not taken, routing becomes unpredictable.

This guideline applies to all Cisco networking devices in a network. If any link in the network has a reduced MTU size, all devices must be changed, not just the devices directly connected to the link.



Note

Do not set the **lsp-mtu** command (network layer) to a value greater than the link MTU size that is set with the **mtu** command (physical layer).

To be certain about a link MTU size, use the [show isis interface](#), on page 863 command to display the value.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to set the MTU size to 1300 bytes:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# isp-mtu 1300
```

Related Commands

| Command | Description |
|--|---|
| mtu | Adjusts the maximum packet size or MTU size. |
| show isis interface, on page 863 | Displays information about the IS-IS interface. |

lsp-password

To configure the link-state packet (LSP) authentication password, use the **lsp-password** command in router configuration mode. To remove the **lsp-password** command from the configuration file and disable link-state packet authentication, use the **no** form of this command.

```
lsp-password [[hmac-md5| text] [clear| encrypted] password] keychain keychain-name [level {1| 2}]
[send-only] [snp send-only]
```

```
no lsp-password [[hmac-md5| text] [clear| encrypted] password] keychain keychain-name [level {1| 2}]
[send-only] [snp send-only]
```

Syntax Description

| | |
|------------------------|---|
| hmac-md5 | Specifies that the password uses HMAC-MD5 authentication. |
| text | Specifies that the password uses clear text password authentication. |
| clear | Specifies that the password be unencrypted. |
| encrypted | Specifies that the password be encrypted using a two-way algorithm. |
| <i>password</i> | Authentication password you assign. |
| keychain | (Optional) Specifies a keychain. |
| <i>keychain-name</i> | Name of the keychain. |
| level { 1 2 } | (Optional) Specifies the password for Level 1 or Level 2 independently. |
| send-only | (Optional) Adds passwords to LSP and sequence number protocol (SNP) data units when they are sent. Does not check for authentication in received LSPs or sequence number PDUs (SNPs). |
| snp send-only | (Optional) Adds passwords to SNP data units when they are sent. Does not check for authentication in received SNPs. This option is available when the text keyword is specified. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a **text** password is configured, it is exchanged as clear text. Therefore, the **lsp-password** command provides limited security.

When an **HMAC-MD5** password is configured, the password is never sent over the network and is instead used to calculate a cryptographic checksum to ensure the integrity of the exchanged data.

The recommended password configuration is that both incoming and outgoing SNPs be authenticated.

**Note**

To disable SNP password checking, the **snp send-only** keywords must be specified in the **lsp-password** command.

To configure an additional password, use the **lsp-password accept** command.

Specify a key chain to enable key chain authentication between two IS-IS peers. Use the **keychain** *keychain-name* keyword and argument to implement hitless key rollover for authentication.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure separate Level 1 and Level 2 LSP and SNP passwords, one with HMAC-MD5 authentication and encryption and one with clear text password authentication and no encryption:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-password hmac-md5 clear password1 level 1
RP/0/RSP0/CPU0:router(config-isis)# lsp-password text clear password2 level 2
```

Related Commands

| Command | Description |
|---|--|
| lsp-password accept , on page 771 | Configures an additional LSP password when one LSP password is already configured for a level. |

lsp-password accept

To configure an additional link-state packet (LSP) authentication password, use the **lsp-password accept** command in router configuration mode. To remove the **lsp-password accept** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

lsp-password accept {clear| encrypted} *password* [level {1| 2}]

no lsp-password accept [{clear| encrypted} *password* [level {1| 2}]]

Syntax Description

| | |
|------------------------|---|
| clear | Specifies that the password be unencrypted. |
| encrypted | Specifies that the password be encrypted using a two-way algorithm. |
| <i>password</i> | Authentication password you assign. |
| level { 1 2 } | (Optional) Specifies the password for Level 1 or Level 2 independently. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **lsp-password accept** command adds an additional password for use when the system validates incoming LSPs and sequence number PDUs (SNPs). An LSP password must be configured using the **lsp-password** command before an accept password can be configured for the corresponding level.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure an Level 1 LSP and SNP password:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-password accept encrypted password1 level 1
```

Related Commands

| Command | Description |
|--|--|
| lsp-password , on page 769 | Configures an authentication LSP password. |

lsp-refresh-interval

To set the time between regeneration of link-state packets (LSPs) that contain different sequence numbers, use the **lsp-refresh-interval** command in router configuration mode. To restore the default refresh interval, use the **no** form of this command.

lsp-refresh-interval *seconds* [**level** {**1** | **2**}]

no lsp-refresh-interval [*seconds* [**level** {**1** | **2**}]]

Syntax Description

| | |
|--------------------------------------|--|
| <i>seconds</i> | Refresh interval (in seconds). Range is 1 to 65535 seconds. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

seconds : 900 seconds (15 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The refresh interval determines the rate at which the software periodically sends the route topology information that it originates. This behavior is done to keep the information from becoming too old. By default, the refresh interval is 900 seconds (15 minutes).

LSPs must be refreshed periodically before their lifetimes expire. The refresh interval must be less than the LSP lifetime specified with this router command. Reducing the refresh interval reduces the amount of time that undetected link-state database corruption can persist at the cost of increased link utilization. (This event is extremely unlikely, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to change the LSP refresh interval to 10,800 seconds (3 hours):

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# lsp-refresh-interval 10800
```

Related Commands

| Command | Description |
|---|--|
| max-lsp-lifetime, on page 778 | Sets the maximum time that LSPs persist without being refreshed. |

maximum-paths (IS-IS)

To configure the maximum number of parallel routes that an IP routing protocol will install in the routing table, use the **maximum-paths** command in address family configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum-paths *maximum*

no maximum-paths

Syntax Description

| | |
|----------------|--|
| <i>maximum</i> | Maximum number of parallel routes that IS-IS can install in a routing table. Range is 1 to 8 |
|----------------|--|

Command Default

maximum : 8 routes

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to allow a maximum of 16 paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# maximum-paths 8
```

maximum-redistributed-prefixes (IS-IS)

To specify an upper limit on the number of redistributed prefixes (subject to summarization) that the Intermediate System-to-Intermediate System (IS-IS) protocol advertises, use the **maximum-redistributed-prefixes** command in address family mode. To disable this feature, use the **no** form of this command.

maximum-redistributed-prefixes *maximum* [level {1| 2}]

no maximum-redistributed-prefixes [*maximum* [level {1| 2}]]

Syntax Description

| | |
|----------------|---|
| <i>maximum</i> | Maximum number of redistributed prefixes advertised. Range is 1 to 28000. |
|----------------|---|

| | |
|------------------------|---|
| level { 1 2 } | (Optional) Specifies maximum prefixes for Level 1 or Level 2. |
|------------------------|---|

Command Default

maximum: 10000

level : 1-2

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum-redistributed-prefixes** command to prevent a misconfiguration from resulting in redistribution of excess prefixes. If IS-IS encounters more than the maximum number of prefixes, it sets a bi-state alarm. If the number of to-be-redistributed prefixes drops back to the maximum or lower—either through reconfiguration or a change in the redistribution source—IS-IS clears the alarm.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to specify the number of redistributed prefixes at 5000 for Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# maximum-redistributed-prefixes 5000 level 2
```

max-lsp-lifetime

To set the maximum time that link-state packets (LSPs) persist without being refreshed, use the **max-lsp-lifetime** command in router configuration mode. To restore the default time, use the **no** form of this command.

max-lsp-lifetime *seconds* [**level** {**1** | **2**}]

no max-lsp-lifetime [*seconds* [**level** {**1** | **2**}]]

Syntax Description

| | |
|--------------------------------------|--|
| <i>seconds</i> | Lifetime (in seconds) of the LSP. Range from 1 to 65535 seconds. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

seconds : 1200 seconds (20 minutes)

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might need to adjust the maximum LSP lifetime if you change the LSP refresh interval with the **lsp-refresh-interval** command. The maximum LSP lifetime must be greater than the LSP refresh interval.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to set the maximum time that the LSP persists to 11,000 seconds (more than 3 hours):

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# max-lsp-lifetime 11000
```

Related Commands

| Command | Description |
|--|--------------------------------|
| lsp-refresh-interval , on page 773 | Sets the LSP refresh interval. |

mesh-group (IS-IS)

To optimize link-state packet (LSP) flooding in highly meshed networks, use the **mesh-group** command in interface configuration mode. To remove a subinterface from a mesh group, use the **no** form of this command.

mesh-group {*number*| **blocked**}

no mesh-group

Syntax Description

| | |
|----------------|--|
| <i>number</i> | Number identifying the mesh group of which this interface is a member. Range is 1 to 4294967295. |
| blocked | Specifies that no LSP flooding takes place on this interface. |

Command Default

There is no mesh group configuration (normal LSP flooding).

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

LSPs first received on subinterfaces that are not part of a mesh group are flooded to all other subinterfaces in the usual way.

LSPs first received on subinterfaces that are part of a mesh group are flooded to all interfaces except those in the same mesh group. If the **blocked** keyword is configured on a subinterface, then a newly received LSP is not flooded out over that interface.

To minimize the possibility of incomplete flooding, you should allow unrestricted flooding over at least a minimal set of links in the mesh. Selecting the smallest set of logical links that covers all physical paths results in very low flooding, but less robustness. Ideally you should select only enough links to ensure that LSP flooding is not detrimental to scaling performance, but enough links to ensure that under most failure scenarios, no router is logically disconnected from the rest of the network. In other words, blocking flooding on all links permits the best scaling performance, but there is no flooding. Permitting flooding on all links results in very poor scaling performance.

**Note**

See RFC 2973 for details about the mesh group specification.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

In the following example, six interfaces are configured in three mesh groups. LSPs received are handled as follows:

- LSPs first received by GigabitEthernet interface 0/1/0/0 are flooded to all interfaces except GigabitEthernet 0/1/0/1 (which is part of the same mesh group) and GigabitEthernet 0/3/0/0 (which is blocked).
- LSPs first received by GigabitEthernet 0/2/0/1 are flooded to all interfaces except GigabitEthernet 0/2/0/0 (which is part of the same mesh group) and GigabitEthernet 0/3/0/0 (which is blocked).
- LSPs first received by GigabitEthernet 0/3/0/0 are not ignored, but flooded as usual to all interfaces.
- LSPs received first through GigabitEthernet 0/3/0/1 are flooded to all interfaces, except GigabitEthernet 0/3/0/0 (which is blocked).

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 10
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 11
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group 12
RP/0/RSP0/CPU0:router(config-isis-if)# exit
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# mesh-group blocked
```

metric (IS-IS)

To configure the metric for an Intermediate System-to-Intermediate System (IS-IS) interface, use the **metric** command in address family or interface address family configuration mode. To restore the default metric value, use the **no** form of this command.

metric [*default-metric*] **maximum** [level {1|2}]

no metric [{*default-metric*] **maximum** [level {1|2}]]

Syntax Description

| | |
|------------------------|---|
| <i>default-metric</i> | Metric assigned to the link and used to calculate the cost from each other router using the links in the network to other destinations. Range is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. Note Setting the default metric under address family results in setting the same metric for all interfaces that is associated with the address family. Setting a metric value under an interface overrides the default metric |
| maximum | Specifies maximum wide metric. All routers exclude this link from their shortest path first (SPF). |
| level { 1 2 } | (Optional) Specifies the SPF calculation for Level 1 or Level 2 independently. |

Command Default

default-metric : Default is 10.

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Interface address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Specifying the **level** keyword resets the metric only for the specified level. We highly recommend that you configure metrics on all interfaces.

Set the default metric under address family to set the same metric for all interfaces that is associated with the address family. Set a metric value under an interface to override the default metric.

We highly recommend that you configure metrics on all interfaces.

Metrics of more than 63 cannot be used with narrow metric style.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure Packet-over-SONET/SDH 0/1/0/1 interface with a default link-state metric cost of 15 for Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# metric 15 level 1
```

The following example shows how to configure a metric cost of 15 for all interfaces under address family IPv4 unicast for level 2:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# metric 15 level 2
```

Related Commands

| Command | Description |
|--|---|
| metric-style narrow, on page 784 | Configures a router running IS-IS so that it generates and accepts old-style TLV objects. |
| metric-style transition, on page 786 | Configures the software to generate and accept both old-style and new-style TLV objects. |
| metric-style wide, on page 788 | Configures the software to generate and accept only new-style TLV objects objects. |

metric-style narrow

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept old-style type, length, and value (TLV) objects, use the **metric-style narrow** command in address family configuration mode. To remove the **metric-style narrow** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style narrow [**transition**] [**level** {**1** | **2**}]

no metric-style narrow [**transition**] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|---|
| transition | (Optional) Instructs the router to generate and accept both old-style and new-style TLV objects. It generates only old-style TLV objects. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

Old-style TLVs are generated.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects with wider metric fields than old-style TLV objects. By default, the router generates old-style TLV objects only. To perform Multiprotocol Label Switching traffic engineering (MPLS TE), a router must generate new-style TLV objects.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the router to generate and accept only old-style TLV objects on router Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style narrow level 1
```

Related Commands

| Command | Description |
|--|--|
| metric-style transition, on page 786 | Configures a router to generate and accept both old-style and new-style TLV objects. |
| metric-style wide, on page 788 | Configures a router to generate and accept only new-style TLV objects. |

metric-style transition

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept both old-style and new-style type, length, and value (TLV) objects, use the **metric-style transition** command in address family configuration mode. To remove the **metric-style transition** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style transition [level {1|2}]

no metric-style transition [level {1|2}]

Syntax Description

| | |
|------------------------|---|
| transition | Instructs the router to generate and accept both old-style and new-style TLV objects. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

Old-style TLVs are generated, if this command is not configured.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects which have wider metric fields than old-style TLV objects. By default, the router generates old-style TLV objects only. To perform Multiprotocol Label Switching traffic engineering (MPLS TE), a router needs to generate new-style TLV objects.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the router to generate and accept both old-style and new-style TLV objects on Level 2:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style transition level 2
```

Related Commands

| Command | Description |
|--|--|
| metric-style narrow, on page 784 | Configures a router to generate and accept only old-style TLV objects. |
| metric-style wide, on page 788 | Configures a router to generate and accept only new-style TLV objects. |

metric-style wide

To configure the Intermediate System-to-Intermediate System (IS-IS) software to generate and accept only new-style type, length, and value (TLV) objects, use the **metric-style wide** command in address family configuration mode. To remove the **metric-style wide** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

metric-style wide [**transition**] [**level** {**1** | **2**}]

no metric-style wide [**transition**] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|---|
| transition | (Optional) Instructs the router to generate and accept both old-style and new-style TLV objects. It generates only new-style TLV objects. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

Old-style TLV lengths are generated, if this command is not configured.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

IS-IS traffic engineering extensions include new-style TLV objects with wider metric fields than old-style TLV objects. If you enter the **metric-style wide** command, a router generates and accepts only new-style TLV objects. Therefore, the router uses less memory and fewer other resources rather than generating both old-style and new-style TLV objects.

To perform MPLS traffic engineering, a router needs to generate new-style TLV objects.



Note

This discussion of metric styles and transition strategies is oriented toward traffic engineering deployment. Other commands and models might be appropriate if the new-style TLV objects are desired for other reasons. For example, a network may require wider metrics, but might not use traffic engineering.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a router to generate and accept only new-style TLV objects on Level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# metric-style wide level 1
```

Related Commands

| Command | Description |
|--|--|
| metric-style narrow, on page 784 | Configures a router to generate and accept only old-style TLV objects. |

min-lsp-arrivaltime

To control the rate of incoming LSPs (link-state packets) LSPs, use the **min-lsp-arrivaltime** command in router configuration mode. To remove this function use the **no** form of this command.

min-lsp-arrivaltime [**initial-wait initial**] [**secondary-wait secondary**] [**maximum-wait maximum**] [**level {1| 2}**]

no min-lsp-arrivaltime [**initial-wait initial**] [**secondary-wait secondary**] [**maximum-wait maximum**] [**level {1| 2}**]

Syntax Description

| | |
|---------------------------------|--|
| initial-wait initial | Initial LSP calculation delay (in milliseconds). Range is 0 to 120000. |
| secondary-wait secondary | Hold time between the first and second LSP calculations (in milliseconds). Range is 0 to 120000. |
| maximum-wait maximum | Maximum interval (in milliseconds) between two consecutive LSP calculations. Range is 0 to 120000. |
| level { 1 2 } | (Optional) Enables the LSP interval configuration for Level 1 or Level 2 independently. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command can be used to protect a router against the possible instability of its neighbor's LSPs.

The command parameters are similar to **lsp-gen-interval** command and neighbors **lsp-gen-interval** values can be used to set the **min-lsp-arrivaltime**



Note

The initial-wait of minimum-lsp-arrival has no use in computing maximum counts and maximum window sizes of the LSP arrival time parameter.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| isis | read, write |

Examples

The following example shows how to configure min-lsp-arrival time commands:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config)# router isis isp min-lsp-arrivaltime  
RP/0/RSP0/CPU0:router(config)# router isis 1 min- lsp-arrivaltime initial-wait  
RP/0/RSP0/CPU0:router(config)#router isis 1 min-lsp-arrivaltime maximum-wait  
RP/0/RSP0/CPU0:router(config)#router isis 1 min-lsp-arrivaltime secondary-wait
```

mpls ldp auto-config

To enable Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) interface auto-configuration, use the **mpls ldp auto-config** command in IPv4 address family configuration mode. To disable LDP IGP auto-configuration, use the **no** form of this command.

mpls ldp auto-config

no mpls ldp auto-config

Syntax Description This command has no keywords or arguments.

Command Default LDP IGP auto-configuration is disabled.

Command Modes IPv4 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mpls ldp auto-config** command to automatically configure LDP on a set of interfaces associated with a specified IGP instance. Further, LDP IGP auto-configuration provides a means to block LDP from being enabled on a specified interface. If you do not want an IS-IS interface to have LDP enabled, use the **igp auto-config disable** command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to enable LDP IGP auto-configuration:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls ldp auto-config
```

Related Commands

| Command | Description |
|--------------------------------|---|
| igp auto-config disable | Disables LDP IGP auto-configuration for a specific interface. |

mpls ldp sync (IS-IS)

To configure Label Distribution Protocol (LDP) IS-IS synchronization, use the **mpls ldp sync** command in interface address family configuration mode. To disable LDP synchronization, use the **no** form of this command.

mpls ldp sync [level {1|2}]

no mpls ldp sync [level {1|2}]

Syntax Description

level { 1 | 2 } (Optional) Sets LDP synchronization for the specified level.

Command Default

If a level is not specified, LDP synchronization is set for both levels.

Command Modes

Interface address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

MPLS VPN traffic forwarded using LDP labels can be dropped in the following instances:

- A new link is introduced in the network and IS-IS has converged before LDP establishes labels.
- An existing LDP session goes down while IS-IS adjacency is intact over the link.

In both instances, outbound LDP labels are not available for forwarding MPLS traffic. LDP IS-IS synchronization addresses the traffic drop. When the **mpls ldp sync** command is configured, IS-IS advertises the maximum possible link metric until LDP has converged over the link. The link is less preferred and least used in forwarding MPLS traffic. When LDP establishes the session and exchanges labels, IS-IS advertises the regular metric over the link.



Note

IS-IS advertises the maximum metric -1 (16777214) if wide metrics are configured since the maximum wide metric is specifically used for link exclusion from the shortest path first algorithm (SPF) (RFC 3784). However, the maximum narrow metric is unaffected by this definition.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to enable LDP IS-IS synchronization:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# mpls ldp sync
```

Related Commands

| Command | Description |
|--|---|
| show isis interface, on page 863 | Displays information about the IS-IS interfaces |

mpls traffic-eng (IS-IS)

To configure a router running the Intermediate System-to-Intermediate System (IS-IS) protocol to flood Multiprotocol Label Switching traffic engineering (MPLS TE) link information into the indicated IS-IS level, use the **mpls traffic-eng** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng {level-1| level-1-2| level-2-only}

no mpls traffic-eng [level-1| level-1-2| level-2-only]

Syntax Description

| | |
|---------------------|-----------------------------------|
| level-1 | Specifies routing level 1. |
| level-1-2 | Specifies routing levels 1 and 2. |
| level-2-only | Specifies routing level 2. |

Command Default

Flooding is disabled.

Command Modes

IPv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mpls traffic-eng** command, which is part of the routing protocol tree, to flood link resource information (such as available bandwidth) for appropriately configured links in the link-state packet (LSP) of the router.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to turn on MPLS traffic engineering for IS-IS level 1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# mpls traffic-eng level-1
```

Related Commands

| Command | Description |
|---|--|
| mpls traffic-eng router-id (IS-IS), on page 802 | Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface. |

mpls traffic-eng multicast-intact (IS-IS)

To enable multicast-intact for Intermediate System-to-Intermediate System (IS-IS) routes with Protocol-Independent Multicast (PIM) and Multiprotocol Label Switching (MPLS) traffic engineering, use the **mpls traffic-eng multicast-intact** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng multicast-intact

no mpls traffic-eng [multicast-intact]

Syntax Description This command has no keywords or arguments.

Command Default Multicast-intact is disabled.

Command Modes IPv4 address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If Multiprotocol Label Switching Traffic Engineering (MPLS-TE) is configured through the IS-IS routing domain and multicast protocols (like Protocol Independent Multicast [PIM]) are also enabled, then use the **mpls traffic-eng multicast-intact** command to install nontraffic engineering next hops in the Routing Information Base (RIB) for use by multicast. The installation of IP-only next hops is in addition to the installation of the standard set of paths for a prefix, which might be through traffic engineered tunnels.

The **mpls traffic-eng multicast-intact** command allows PIM to use the native hop-by-hop neighbors even though the unicast routing is using MPLS TE tunnels.

Examples The following example shows how to enable the multicast-intact feature:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls traffic-eng multicast-intact
```

Related Commands

| Command | Description |
|---|--|
| show isis route, on page 887 | Displays IP reachability information for an IS-IS instance, optionally for multicast-intact. |
| show isis topology, on page 902 | Displays a list of connected IS-IS routers in all areas, optionally for multicast-intact. |

mpls traffic-eng path-selection ignore overload

To ensure that label switched paths (LSPs) are not disabled when routers have the Intermediate System-to-Intermediate System (IS-IS) overload bit set, use the **mpls traffic-eng path-selection ignore overload** command in global configuration mode. To disable this override, use the **no** form of this command.

mpls traffic-eng path-selection ignore overload

no mpls traffic-eng path-selection ignore overload

Command Default No default behavior or values

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the IS-IS overload bit avoidance feature is activated, which means that they are still available for use label switched paths (LSPs), all nodes with the overload bit set, including the following nodes, are ignored:

- head nodes
- mid nodes
- tail nodes

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | mpls-te | read, write |

Examples The following example shows how to activate IS-IS overload bit avoidance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# mpls traffic-eng path-selection ignore overload
```

The following example shows how to deactivate IS-IS overload bit avoidance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# no mpls traffic-eng path-selection ignore overload
```

Related Commands

| Command | Description |
|---|---|
| set-overload-bit, on page 831 | Configures a router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations. |

mpls traffic-eng router-id (IS-IS)

To specify the Multiprotocol Label Switching traffic engineering (MPLS TE) router identifier for the node, use the **mpls traffic-eng router-id** command in IPv4 address family configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id *{ip-address| type interface-path-id}*

no mpls traffic-eng [router-id]

Syntax Description

| | |
|--------------------------|---|
| <i>ip-address</i> | IP address in four-part, dotted-decimal notation. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

Global router identifier is used.

Command Modes

IPv4 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router ID of the destination node, because that is the address used by the traffic engineering topology database at the tunnel head for its path calculation.



Note

We recommend that loopback interfaces be used for MPLS TE, because they are more stable than physical interfaces.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# mpls traffic-eng router-id Loopback0
```

Related Commands

| Command | Description |
|---|--|
| mpls traffic-eng (IS-IS), on page 796 | Turns on flooding of MPLS traffic engineering link information in the indicated IGP level or area. |

net

To configure an Intermediate System-to-Intermediate System (IS-IS) network entity title (NET) for the routing instance, use the **net** command in router configuration mode. To remove the **net** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

net *network-entity-title*

no net *network-entity-title*

Syntax Description

| | |
|----------------------|--|
| network-entity-title | NET that specifies the area address and the system ID for an ISIS routing process. |
|----------------------|--|

Command Default

No NET is configured. The IS-IS instance is not operational, because a NET is mandatory.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Under most circumstances, one and only one NET should be configured.

A NET is a network service access point (NSAP) where the last byte is always 0. On a Cisco router running IS-IS, a NET can be 8 to 20 bytes in length. The last byte is always the n-selector and must be 0. The n-selector indicates to which transport entity the packet is sent. An n-selector of 0 indicates no transport entity and means that the packet is for the routing software of the system.

The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2).

All bytes preceding the system ID are the area ID.

A maximum of three NETs for each router is allowed. In rare circumstances, it is possible to configure two or three NETs. In such a case, the area this router is in has three area addresses. Only one area still exists, but it has more area addresses.

Configuring multiple NETs can be temporarily useful in network reconfiguration in which multiple areas are merged, or in which one area is split into more areas. Multiple area addresses enable you to renumber an area individually as needed.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure a router with NET area ID 47.0004.004d.0001 and system ID 0001.0c11.1110:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 47.0004.004d.0001.0001.0c11.1110.00
```

Related Commands

| Command | Description |
|--|--|
| log adjacency changes (IS-IS), on page 758 | Configures the routing level for an instance of the IS-IS routing process. |
| router isis, on page 827 | Enables the IS-IS routing protocol and specifies an IS-IS instance. |

nsf (IS-IS)

To enable nonstop forwarding (NSF) on the next restart, use the **nsf** command in router configuration mode. To restore the default setting, use the **no** form of this command.

nsf {cisco|ietf}

no nsf {cisco|ietf}

Syntax Description

| | |
|--------------|---|
| cisco | Specifies Cisco-proprietary NSF restart. |
| ietf | Specifies Internet Engineering Task Force (IETF) NSF restart. |

Command Default

NSF is disabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

NSF allows an Intermediate System-to-Intermediate System (IS-IS) instance to restart using checkpointed adjacency and link-state packet (LSP) information, and to perform restart with no impact on its neighbor routers. In other words, there is no impact on other routers in the network due to the destruction and recreation of adjacencies and the system LSP.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to enable Cisco proprietary NSF:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf cisco
```

Related Commands

| Command | Description |
|---|--|
| nsf interface-expires , on page 808 | Configures the number of resends of an acknowledged NSF-restart acknowledgment. |
| nsf interface-timer , on page 810 | Configures the time interval after which an unacknowledged IETF NSF restart attempt is repeated. |
| nsf lifetime (IS-IS) , on page 812 | Configures the maximum route lifetime following an NSF restart. |

nsf interface-expires

To configure the number of resends of an acknowledged nonstop forwarding (NSF)-restart acknowledgment, use the **nsf interface-expires** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf interface-expires *number*

no nsf interface-expires

Syntax Description

| | |
|--------|-------------------------------------|
| number | Number of resends. Range is 1 to 3. |
|--------|-------------------------------------|

Command Default

number : 3 resends

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a hello packet sent with the NSF restart flag set is not acknowledged, it is re-sent. Use the **nsf interface-expires** command to control the number of times the NSF hello is re-sent. When this limit is reached on an interface, any neighbor previously known on that interface is assumed to be down and the initial shortest path first (SPF) calculation is permitted, provided that all other necessary conditions are met.

The total time period available for adjacency reestablishment (`interface-timer * interface-expires`) should be greater than the expected total NSF restart time.

The **nsf interface-expires** command applies only to Internet Engineering Task Force (IETF)-style NSF. It has no effect if Cisco-proprietary NSF is configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to allow only one retry attempt on each interface if an IETF NSF restart signal is not acknowledged:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# nsf ietf  
RP/0/RSP0/CPU0:router(config-isis)# nsf interface-expires 1
```

Related Commands

| Command | Description |
|--|--|
| hello-multiplier, on page 739 | Specifies the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down. |
| nsf interface-timer, on page 810 | Configures the time interval after which an unacknowledged IETF NSF restart attempt is repeated. |

nsf interface-timer

To configure the time interval after which an unacknowledged Internet Engineering Task Force (IETF) nonstop forwarding (NSF) restart attempt is repeated, use the **nsf interface-timer** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf interface-timer *seconds*

no nsf interface-timer

| Syntax Description | seconds | NSF restart time interval (in seconds). Range is 3 to 20 seconds. |
|--------------------|---------|---|
|--------------------|---------|---|

| Command Default | <i>seconds</i> : 10 seconds |
|-----------------|-----------------------------|
|-----------------|-----------------------------|

| Command Modes | Router configuration |
|---------------|----------------------|
|---------------|----------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the IETF NSF restart process begins, hello packets send an NSF restart flag that must be acknowledged by the neighbors of the router. Use the **nsf interface-timer** command to control the restart time interval after the hello packet is re-sent. The restart time interval need not match the hello interval.

The **nsf interface-timer** command applies only to IETF-style NSF. It has no effect if Cisco proprietary NSF is configured.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples

The following example shows how to ensure that a hello packet with the NSF restart flag set is sent again every 5 seconds until the flag is acknowledged:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf ietf
```

```
RP/0/RSP0/CPU0:router(config-isis)# nsf interface-timer 5
```

Related Commands

| Command | Description |
|---|---|
| nsf interface-expires, on page 808 | Configures the number of resends of an acknowledged NSF-restart acknowledgment. |
| hello-interval (IS-IS), on page 737 | Specifies the length of time between hello packets that the software sends. |

nsf lifetime (IS-IS)

To configure the maximum route lifetime following a nonstop forwarding (NSF) restart, use the **nsf lifetime** command in router configuration mode. To restore the default value, use the **no** form of this command.

nsf lifetime *seconds*

no nsf lifetime

Syntax Description

| | |
|---------|--|
| seconds | Maximum route lifetime (in seconds) following an NSF restart. Range is 5 to 300 seconds. |
|---------|--|

Command Default

seconds : 60 seconds (1 minute)

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **nsf lifetime** command to set the maximum available time for the reacquisition of checkpointed adjacencies and link-state packets (LSPs) during a Cisco proprietary NSF restart. LSPs and adjacencies not recovered during this time period are abandoned, thus causing changes to the network topology.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the router to allow only 20 seconds for the entire NSF process:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# nsf cisco
RP/0/RSP0/CPU0:router(config-isis)# nsf lifetime 20
```

passive (IS-IS)

To suppress Intermediate System-to-Intermediate System (IS-IS) packets from being transmitted to the interface and received packets from being processed on the interface, use the **passive** command in interface configuration mode. To restore IS-IS packets coming to an interface, use the **no** form of this command.

passive

no passive

Command Default Interface is active.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to configure the router to suppress IS-IS packets on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# passive
```

| Related Commands | Command | Description |
|------------------|---|--|
| | suppressed, on page 919 | Allows the IS-IS interface to participate in forming adjacencies without advertising connected prefixes in the LSPs. |

point-to-point

To configure a network of only two networking devices that use broadcast media and the integrated Intermediate System-to-Intermediate System (IS-IS) routing protocol to function as a point-to-point link instead of a broadcast link, use the **point-to-point** command in interface configuration mode. To disable the point-to-point usage, use the **no** form of this command.

point-to-point

no point-to-point

Syntax Description This command has no keywords or arguments.

Command Default Interface is treated as broadcast if connected to broadcast media.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **point-to-point** command only on broadcast media in a network with two networking devices. The command causes the system to issue packets point-to-point rather than as broadcasts. Configure the command on both networking devices in the network.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to configure a 10-Gb Ethernet interface to act as a point-to-point interface:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE 0/6/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# point-to-point
```

priority (IS-IS)

To configure the priority of designated routers, use the **priority** command in interface configuration mode. To reset the default priority, use the **no** form of this command.

priority *value* [**level** {**1** | **2**}]

no priority [*value*] [**level** {**1** | **2**}]

Syntax Description

| | |
|--------------------------------------|--|
| <i>value</i> | Priority of a router. Range is 0 to 127. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

value : 64

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Priorities can be configured for Level 1 and Level 2 independently. Specifying Level 1 or Level 2 resets priority only for Level 1 or Level 2 routing, respectively. Specifying no level allows you to configure all levels.

The priority is used to determine which router on a LAN is the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS.

In the Intermediate System-to-Intermediate System (IS-IS) protocol, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes online, it takes over the role from the current DIS. For equal priorities, the higher MAC address breaks the tie.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to give Level 1 routing priority by setting the priority level to 80. This router is now more likely to become the DIS.

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# interface TenGigE0/6/0/0  
RP/0/RSP0/CPU0:router(config-isis-if)# priority 80 level 1
```

propagate level

To propagate routes from one Intermediate System-to-Intermediate System (IS-IS) level into another level, use the **propagate level** command in address family configuration mode. To disable propagation, use the **no** form of this command.

propagate level {1|2} **into level** {1|2} **route-policy** *route-policy-name*

no propagate level {1|2} **into level** {1|2}

Syntax Description

| | |
|--|---|
| level { 1 2 } | Propagates from routing Level 1 or Level 2 routes. |
| into | Propagates from Level 1 or Level 2 routes into Level 1 or Level 2 routes. |
| route-policy <i>route-policy-name</i> | Specifies a configured route policy. |

Command Default

Route leaking (Level 2 to Level 1) is disabled.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In general, route propagation from Level 1 to Level 2 is automatic. You might want to use this command to better control which Level 1 routes can be propagated into Level 2.

Propagating Level 2 routes into Level 1 is called *route leaking*. Route leaking is disabled by default. That is, Level 2 routes are not automatically included in Level 1 link-state packets (LSPs). If you want to leak Level 2 routes into Level 1, you must enable that behavior by using this command.

Propagation from Level 1 into Level 1 and from Level 2 into Level 2 is not allowed.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to redistribute Level 2 routes to Level 1:

```
RP/0/RSP0/CPU0:router(config)# ipv4 access-list 101 permit ip 10.0.0.0 255.0.0.0 10.1.0.1
0.255.255.255
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.1234.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# propagate level 2 into level 1 route-policy policy_a
```

Related Commands

| Command | Description |
|---|---|
| redistribute (IS-IS), on page 819 | Redistributes routes from one routing domain into a specified IS-IS instance. |

redistribute (IS-IS)

To redistribute routes from one routing protocol into Intermediate System-to-Intermediate System (IS-IS), use the **redistribute** command in address family configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)

```
redistribute bgp process-id [level-1 | level-2 | level-1-2] [metric metric-value] [metric-type {internal | external}] [route-policy route-policy-name]
```

```
no redistribute
```

Connected Routes

```
redistribute connected [level-1 | level-2 | level-1-2] [metric metric-value] [metric-type {internal | external}] [route-policy route-policy-name]
```

```
no redistribute
```

Intermediate System-to-Intermediate System (IS-IS)

```
redistribute isis process-id [level-1 | level-2 | level-1-2] [metric metric-value] [metric-type {internal | external}] [route-policy route-policy-name]
```

```
no redistribute
```

Open Shortest Path First (OSPF)

```
redistribute ospf process-id [level-1 | level-2 | level-1-2] [match {external [1 | 2] | internal | nssa-external [1 | 2]}] [metric metric-value] [metric-type {internal | external}] [route-policy route-policy-name]
```

```
no redistribute
```

Open Shortest Path First Version 3 (OSPFv3)

```
redistribute ospfv3 process-id [level-1 | level-2 | level-1-2] [match {external [1 | 2] | internal | nssa-external [1 | 2]}] [metric metric-value] [metric-type {internal | external}] [route-policy route-policy-name]
```

```
no redistribute
```

Static Routes

```
redistribute static [level-1 | level-2 | level-1-2] [metric metric-value] [metric-type {1 | 2}] [route-policy route-policy-name]
```

```
no redistribute
```

Syntax Description

| | |
|---|--|
| <i>process-id</i> | <p>For the bgp keyword, an autonomous system number has the following ranges:</p> <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. <p>For the isis keyword, an IS-IS instance identifier from which routes are to be redistributed.</p> <p>For the ospf keyword, an OSPF process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> <p>For the ospfv3 keyword, an OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| level-1 | (Optional) Specifies that redistributed routes are advertised in the Level-1 LSP of the router. |
| level-1-2 | (Optional) Specifies that redistributed routes are advertised in the Level-1-2 LSP of the router. |
| level-2 | (Optional) Specifies that redistributed routes are advertised in the Level-2 LSP of the router. |
| metric <i>metric-value</i> | (Optional) Specifies the metric used for the redistributed route. Range is 0 to 16777215. The <i>metric-value</i> must be consistent with the IS-IS metric style of the area and topology into which the routes are being redistributed. |
| metric-type { internal external } | <p>(Optional) Specifies the external link type associated with the route advertised into the ISIS routing domain. It can be one of two values:</p> <ul style="list-style-type: none"> • external • internal –Use the internal keyword to set IS-IS internal metric-type • external –Use the external keyword to set IS-IS external metric-type <p>Any route with an internal metric (however large the metric is) is preferred over a route with external metric (however small the metric is).</p> |
| route-policy <i>route-policy-name</i> | (Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to IS-IS. |

match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**] }

(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:

- **internal** —Routes that are internal to a specific autonomous system (intra- and interarea OSPF routes).
- **external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes.
- **nssa-external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes.

For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.

Command Default

Level 2 is configured if no level is specified.

metric-type: internal

match : If no match keyword is specified, all OSPF routes are redistributed.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

When redistributing routes (into IS-IS) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Use the **redistribute** command to control the redistribution of routes between separate IS-IS instances. To control the propagation of routes between the levels of a single IS-IS instance, use the [propagate level](#), on [page 817](#) command.

Only IPv4 OSPF addresses can be redistributed into IS-IS IPv4 address families.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

In this example, IS-IS instance `isp_A` readadvertises all of the routes of IS-IS instance `isp_B` in Level 2 LSP. Note that the **level-2** keyword affects which levels instance `isp_A` advertises the routes in and has no impact on which routes from instance `isp_B` are advertised. (Any Level 1 routes from IS-IS instance `isp_B` are included in the redistribution.)

```
RP/0/RSP0/CPU0:router(config)# router isis isp_A
RP/0/RSP0/CPU0:router(config-isis)# net 49.1234.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute isis isp_B level-2
!
RP/0/RSP0/CPU0:router(config)# router isis isp_B
RP/0/RSP0/CPU0:router(config-isis)# is-type level 1
RP/0/RSP0/CPU0:router(config-isis)# net 49.4567.2222.2222.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
```

Related Commands

| Command | Description |
|--|--|
| propagate level, on page 817 | Propagates routes from one IS-IS level into another level. |

retransmit-interval (IS-IS)

To configure the amount of time between retransmission of each Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) on a point-to-point link, use the **retransmit-interval** command in interface configuration mode. To restore the default value, use the **no** form of this command.

retransmit-interval *seconds* [**level** {**1** | **2**}]

no retransmit-interval [*seconds* [**level** {**1** | **2**}]]

Syntax Description

| | |
|--------------------------------------|--|
| <i>seconds</i> | Time (in seconds) between consecutive retransmissions of each LSP. It is an integer that should be greater than the expected round-trip delay between any two networking devices on the attached network. Range is 0 to 65535 seconds. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

seconds : 5 seconds

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **retransmit-interval** command has no effect on LAN (multipoint) interfaces. On point-to-point links, the value can be increased to enhance network stability.

Because retransmissions occur only when LSPs are dropped, setting this command to a higher value has little effect on convergence. The more neighbors networking devices have, and the more paths over which LSPs can be flooded, the higher this value can be made.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure GigabitEthernet interface 0/2/0/1 for retransmission of IS-IS LSPs every 60 seconds for a large serial line:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# retransmit-interval 60
```

Related Commands

| Command | Description |
|---|--|
| retransmit-throttle-interval, on page 825 | Configures the amount of time between retransmissions of any IS-IS LSPs on a point-to-point interface. |

retransmit-throttle-interval

To configure minimum interval between retransmissions of different Intermediate System-to-Intermediate System (IS-IS) link-state packets (LSPs) on a point-to-point interface, use the **retransmit-throttle-interval** command in interface configuration mode. To remove the command from the configuration file and restore the system to its default condition, use the **no** form of this command.

retransmit-throttle-interval *milliseconds* [**level** {**1** | **2**}]

no retransmit-throttle-interval [*milliseconds* [**level** {**1** | **2**}]]

Syntax Description

| | |
|--------------------------------------|--|
| <i>milliseconds</i> | Minimum delay (in milliseconds) between LSP retransmissions on the interface. Range is 0 to 65535. |
| level { 1 2 } | (Optional) Specifies routing Level 1 or Level 2 independently. |

Command Default

Default is 0.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **retransmit-throttle-interval** command to define the minimum period of time that must elapse between retransmitting any two consecutive LSPs on an interface. The **retransmit-throttle-interval** command may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This command controls the rate at which LSPs can be re-sent on the interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure GigabitEthernet interface 0/2/0/1 to limit the rate of LSP retransmissions to one every 300 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# retransmit-throttle-interval 300
```

Related Commands

| Command | Description |
|---|--|
| lsp-gen-interval , on page 763 | Configures the minimum interval time between regenerating the same LSP. |
| retransmit-interval (IS-IS) , on page 823 | Configures the amount of time between retransmission of each IS-IS LSP over a point-to-point link. |

router isis

To enable the Intermediate System-to-Intermediate System (IS-IS) routing protocol and to specify an IS-IS instance, use the **router isis** command in global configuration mode. To disable IS-IS routing, use the **no** form of this command.

router isis *instance-id*

no router isis *instance-id*

Syntax Description

| | |
|-------------|--|
| instance-id | Name of the routing process. Maximum number of characters is 40. |
|-------------|--|

Command Default

An IS-IS routing protocol is not enabled.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **router isis** command to create an IS-IS routing process. An appropriate network entity title (NET) must be configured to specify the address of the area (Level 1) and system ID of the router. Routing must be enabled on one or more interfaces before adjacencies may be established and dynamic routing is possible.

Multiple IS-IS processes can be configured. Up to eight processes are configurable. A maximum of five IS-IS instances on a system are supported.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure IS-IS for IP routing:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# net 49.0001.0000.0001.00
```

Related Commands

| Command | Description |
|---|--|
| net , on page 804 | Configures an IS-IS NET for the routing process. |

set-attached-bit

To configure an Intermediate System-to-Intermediate System (IS-IS) instance with an attached bit in the Level 1 link-state packet (LSP), use the **set-attached-bit** command in address family configuration mode. To remove the **set-attached-bit** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

set-attached-bit

no set-attached-bit

Command Default Attached bit is not set in the LSP.

Command Modes Address family configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set-attached bit** command to set an IS-IS instance with an attached bit in the Level 1 LSP that allows another IS-IS instance to redistribute Level 2 topology. The attached bit is used when the Level 2 connectivity from another IS-IS instance is advertised by the Level 1 attached bit.

Cisco IOS XR software does not support multiple Level 1 areas in a single IS-IS routing instance. But the equivalent functionality is achieved by redistribution of routes between two IS-IS instances by using the [redistribute \(IS-IS\), on page 819](#) command.

The attached bit is configured for a specific address family only if the **single-topology** command is not configured.



Note If connectivity for the Level 2 instance is lost, the attached bit in the Level 1 instance LSP continues sending traffic to the Level 2 instance and causes the traffic to be dropped.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples

The following example shows how to set the attached bit for a Level 1 instance that allows the Level 2 instance to redistribute routes from the Level 1 instance:

```
RP/0/RSP0/CPU0:router(config)# router isis 1
RP/0/RSP0/CPU0:router(config-isis)# net 49.0001.0001.0001.0001.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute isis 2 level 2
!
RP/0/RSP0/CPU0:router(config-isis-af)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast
!
!
RP/0/RSP0/CPU0:router(config)# router isis 2
RP/0/RSP0/CPU0:router(config-isis)# is-type level-1
RP/0/RSP0/CPU0:router(config-isis)# net 49.0002.0001.0001.0002.00
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# set-attached-bit
!
RP/0/RSP0/CPU0:router(config-isis-af)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-isis-af-if)# address-family ipv4 unicast
```

Related Commands

| Command | Description |
|---|---|
| redistribute (IS-IS), on page 819 | Redistributes routes from one IS-IS instance into another instance. |
| single-topology, on page 910 | Configures the link topology for IPv4 when IPv6 is configured. |

set-overload-bit

To configure the router to signal other routers not to use it as an intermediate hop in their shortest path first (SPF) calculations, use the **set-overload-bit** command in router configuration mode. To remove the designation, use the **no** form of this command.

```
set-overload-bit [on-startup {delay| wait-for-bgp}] [level {1| 2}] [advertise {external| interlevel}]
no set-overload-bit [on-startup {delay| wait-for-bgp}] [level {1| 2}]
```

Syntax Description

| | |
|------------------------|--|
| on-startup | (Optional) Sets the overload bit only temporarily after reboot. |
| <i>delay</i> | (Optional) Time (in seconds) to advertise when the router is overloaded after reboot. Range is 5 to 86400 seconds (86400 seconds = 1 day). |
| wait-for-bgp | (Optional) Sets the overload bit on startup until the Border Gateway Protocol (BGP) signals converge or time out. |
| level { 1 2 } | (Optional) Specifies the overload bit for Level 1 or Level 2 independently. |

Command Default

The overload bit is not set.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set-overload-bit** command to force the router to set the overload bit in its nonpseudonode link-state packets (LSPs). Normally the setting of the overload bit is allowed only when a router experiences problems. For example, when a router is experiencing a memory shortage, the reason might be that the link-state database is not complete, resulting in an incomplete or inaccurate routing table. If the overload bit is set in the LSPs of the unreliable router, other routers can ignore the router in their SPF calculations until it has recovered from its problems. The result is that no paths through the unreliable router are seen by other routers in the Intermediate System-to-Intermediate System (IS-IS) area. However, IP prefixes directly connected to this router are still reachable.

The **set-overload-bit** command can be useful when you want to connect a router to an IS-IS network, but do not want real traffic flowing through it under any circumstances.

Routers with overload bit set are:

- A test router in the lab, connected to a production network.
- A router configured as an LSP flooding server, for example, on a nonbroadcast multiaccess (NBMA) network, in combination with the mesh group feature.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to configure the overload bit:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# set-overload-bit
```

show isis

The **show isis** command displays general information about an IS-IS instance and protocol operation. If the instance ID is not specified, the command shows information about all IS-IS instances.

show isis [**instance** *instance-id*]

Syntax Description

instance *instance-id* (Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only.

Note The instance-id argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For each instance, the first line of output lists the IS-IS instance ID with the following lines identifying the IS-IS system ID, supported levels (level 1, level 2, or level-1-2), configured area addresses, active area addresses, status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF), and the mode in which the last IS-IS process startup occurred.

Next, the status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not). Then redistributed protocols are listed, followed by the administrative distance applied to the redistributed routes.

Finally, the running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface is listed.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis** command:

```
RP/0/RSP0/CPU0:router# show isis
Wed Aug 20 23:54:55.043 PST DST

IS-IS Router: lab
System Id: 0000.0000.0002
IS Levels: level-2-only
Manual area address(es):
 49.1122
Routing for area address(es):
 49.1122
Non-stop forwarding: Disabled
Most recent startup mode: Cold Restart
Topologies supported by IS-IS:
 IPv4 Unicast
  Level-2
   Metric style (generate/accept): Narrow/Narrow
   Metric: 10
   ISPF status: Disabled
   No protocols redistributed
   Distance: 115
Interfaces supported by IS-IS:
 Loopback0 is running passively (passive in configuration)
 POS0/1/0/2 is running actively (active in configuration)
 POS0/1/0/3 is running actively (active in configuration)
```

This table describes the significant fields shown in the display.

Table 57: show isis Field Descriptions

| Field | Description |
|-------------------------------|---|
| IS-IS Router | IS-IS instance ID. |
| System Id | IS-IS system ID. |
| IS Levels | Supported levels for the instance. |
| Manual area address(es) | Domain and area. |
| Routing for area address(es): | Configured area addresses and active area addresses. |
| Non-stop forwarding | Status (enabled or not) and type (Cisco or IETF) of nonstop forwarding (NSF). |
| Most recent startup mode | The mode in which the last IS-IS process startup occurred. |
| Topologies supported by IS-IS | The summary of the status of each configured address family (or just IPv4 unicast if none are configured). |
| Redistributed protocols | List of redistributed protocols, followed by the administrative distance applied to the redistributed routes. |

| Field | Description |
|--------------------------------|---|
| Metric style (generate/accept) | The status of each configured address family (or just IPv4 unicast if none are configured) is summarized. For each level (level 1 or level 2), the metric style (narrow or wide) generated and accepted is listed along with the status of incremental shortest path first (iSPF) computation (enabled or not). |
| Interfaces supported by IS-IS | The running state (active, passive, or disabled) and configuration state (active or disabled) of each IS-IS interface. |

show isis adjacency

To display Intermediate System-to-Intermediate System (IS-IS) adjacencies, use the **show isis adjacency** command in EXEC mode.

show isis [*instance instance-id*] **adjacency** [*level {1|2}*] [*type interface-path-id*] [**detail**] [*systemid system-id*]

Syntax Description

| | |
|------------------------------------|--|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| level { 1 2 } | (Optional) Displays the IS-IS adjacencies for Level 1 or Level 2 independently. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function.</p> |
| detail | (Optional) Displays neighbor IP addresses and active topologies. |
| systemid <i>system-id</i> | (Optional) Displays the information for the specified router only. |

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances. Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis adjacency** command:

```
RP/0/RSP0/CPU0:router# show isis adjacency

IS-IS p Level-1 adjacencies:
System Id      Interface      SNPA          State Hold   Changed NSF      BFD
12a4           PO0/1/0/1     *PtoP*        Up    23     00:00:06 Capable  Init
12a4           Gi0/6/0/2     0004.2893.f2f6 Up    56     00:04:01 Capable  Up

Total adjacency count: 2

IS-IS p Level-2 adjacencies:
System Id      Interface      SNPA          State Hold   Changed NSF      BFD
12a4           PO0/1/0/1     *PtoP*        Up    23     00:00:06 Capable  None
12a4           Gi0/6/0/2     0004.2893.f2f6 Up    26     00:00:13 Capable  Init

Total adjacency count: 2
```

This table describes the significant fields shown in the display.

Table 58: show isis adjacency Field Descriptions

| Field | Description |
|-----------|--|
| Level-1 | Level 1 adjacencies. |
| Level-2 | Level 2 adjacencies. |
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or the hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| Interface | Interface used to reach the neighbor. |
| SNPA | Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor. |
| State | Adjacency state of the neighboring interface. Valid states are Down, Init, and Up. |
| Holdtime | Hold time of the neighbor. |
| Changed | Time the neighbor has been up (in hours:minutes:seconds). |

| Field | Description |
|-------|--|
| NSF | Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism. |
| BFD | Specifies the Bidirectional Forwarding Detection (BFD) status for the interface. Valid status are <ul style="list-style-type: none"> • None—BFD is not configured. • Init—BFD session is not up. One reason is that other side is not yet enabled. • Up—BFD session has been established. • Down—BFD session holdtime expired. |

Related Commands

| Command | Description |
|--|---|
| show isis neighbors, on page 880 | Displays information about IS-IS neighbors. |

show isis adjacency-log

To display the Intermediate System-to-Intermediate System (IS-IS) adjacency log, use the **show isis adjacency-log** command in EXEC mode.

```
show isis adjacency-log [level {1| 2}] [last number| first number]
```

Syntax Description

| | |
|----------------------------|--|
| level { 1 2 } | (Optional) Displays the IS-IS adjacency log for Level 1 or Level 2 independently. |
| last <i>number</i> | (Optional) Specifies that the output is restricted to the last <i>number</i> of entries. Range is 1 to 100. |
| first <i>number</i> | (Optional) Specifies that the output is restricted to the first <i>number</i> of entries. Range is 1 to 100. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis adjacency-log** command:

```
RP/0/RSP0/CPU0:router# show isis adjacency-log

IS-IS 10 Level 1 Adjacency log
When      System      Interface      State  Details
4d00h    12a1          PO0/5/0/0     d -> i
4d00h    12a1          PO0/5/0/0     i -> u  New adjacency
```

show isis adjacency-log

```

4d00h          12a1          Gi0/6/0/0          d -> u          IPv4 Unicast Up
4d00h          12a1          Gi0/6/0/0          u -> d          New adjacency
down
3d17h          12a1          Gi0/6/0/0          d -> u          Interface state
3d17h          12a1          Gi0/6/0/0          u -> d          New adjacency
down
01:44:07      12a1          Gi0/6/0/0          d -> u          Interface state
New adjacency

IS-IS 10 Level 2 Adjacency log
When          System          Interface          State          Details
4d00h          12a1          PO0/5/0/0          d -> i
4d00h          12a1          PO0/5/0/0          i -> u          New adjacency
4d00h          12a1          Gi0/6/0/0          d -> u          IPv4 Unicast Up
4d00h          12a1          Gi0/6/0/0          u -> d          New adjacency
down
3d17h          12a1          Gi0/6/0/0          d -> u          Interface state
3d17h          12a1          Gi0/6/0/0          u -> d          New adjacency
down
01:44:07      12a1          Gi0/6/0/0          d -> u          Interface state
New adjacency

```

This table describes the significant fields shown in the display.

Table 59: show isis adjacency-log Field Descriptions

| Field | Description |
|-----------|--|
| When | Elapsed time (in hh:mm:ss) since the event was logged. |
| System | System ID of the adjacent router. |
| Interface | Specific interface involved in the adjacency change. |
| State | State transition for the logged event. |
| Details | Description of the adjacency change. |

show isis checkpoint adjacency

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint adjacency database, use the **show isis checkpoint adjacency** command in EXEC mode.

show isis [*instance instance-id*] **checkpoint adjacency**

Syntax Description

instance *instance-id* (Optional) Displays the IS-IS checkpoint adjacencies for the specified IS-IS instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint adjacencies for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis checkpoint adjacency** command to display the checkpointed adjacencies. With this information you can restore the adjacency database during a Cisco proprietary nonstop forwarding (NSF) restart. This command, with the **show isis adjacency** command, can be used to verify the consistency of the two databases.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis checkpoint adjacency** command:

```
RP/0/RSP0/CPU0:router# show
isis
checkpoint
```

```

adjacency
Interface      Level  System ID      State  Circuit ID      Chkpt ID
Gi3/0/0/1     1      router-gsr8     Up     0001.0000.0008.04  80011fec
Gi0/4/0/1     1      router-gsr9     Up     0001.0000.0006.01  80011fd8
Gi/0/0/1      2      router-gsr8     Up     0001.0000.0008.04  80011fc4

```

This table describes the significant fields shown in the display.

Table 60: show isis checkpoint adjacency Field Descriptions

| Field | Description |
|------------|--|
| Interface | Interface used to reach the neighbor. |
| Level | Lists either routers with Level 1 or Level 2 adjacency configured. |
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| State | State of the neighboring interface. |
| Circuit ID | Unique ID issued to a circuit at its creation. |
| Chkpt ID | Unique ID issued to the checkpoint at its creation. |

Related Commands

| Command | Description |
|---|---|
| show isis adjacency, on page 836 | Displays IS-IS adjacencies. |
| show isis checkpoint lsp, on page 845 | Displays the IS-IS checkpoint LSP database. |

show isis checkpoint interface

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint interfaces, use the **show isis checkpoint interface** command in EXEC mode.

show isis checkpoint interface

This command has no keywords or arguments.

Command Default No default behavior or values

Command Modes EXEC/EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | isis | read |

Examples The following is sample output from the **show isis checkpoint interface** command:

```
RP/0/RSP0/CPU0:router# show isis checkpoint interface

IS-IS 10 checkpoint interface
Interface      Index  CircNum  DIS Areas  Chkpt ID
PO0/5/0/0      0      0        NONE       80002fe8
Gi0/6/0/0      1      3        L1L2       80002fd0
```

This table describes the significant fields shown in the display.

Table 61: show isis checkpoint interface Field Descriptions

| Field | Description |
|-----------|---------------------------------------|
| Interface | Interface used to reach the neighbor. |

| Field | Description |
|-----------|---|
| Index | Interface index assigned to an interface upon its creation. |
| CircNum | Unique ID issued to a circuit internally. |
| DIS Areas | Designated Intermediate System area. |
| Chkpt ID | Unique ID issued to the checkpoint at its creation. |

show isis checkpoint lsp

To display the Intermediate System-to-Intermediate System (IS-IS) checkpoint link-state packet (LSP) protocol data unit (PDU) identifier database, use the **show isis checkpoint lsp** command in EXEC mode.

show isis [*instance instance-id*] **checkpoint lsp**

Syntax Description

instance *instance-id* (Optional) Displays the IS-IS checkpoint LSPs for the specified instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS checkpoint LSPs for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The checkpointed LSPs displayed by this command are used to restore the LSP database during a Cisco-proprietary nonstop forwarding (NSF) restart. The **show isis checkpoint lsp** command, with the **show isis database** command, may be used to verify the consistency of the two databases.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis checkpoint lsp** command:

```
RP/0/RSP0/CPU0:router#
    show isis checkpoint lsp

Level  LSPID                Chkpt ID
 1     router-gsr6.00-00    80011f9c
```

show isis checkpoint lsp

```

1      router-gsr6.01-00  80011f88
1      router-gsr8.00-00  80011f74
1      router-gsr9.00-00  80011f60
2      router-gsr6.00-00  80011f4c
2      router-gsr6.01-00  80011f38
2      router-gsr8.00-00  80011f24
2      router-gsr9.00-00  80011f10
Total LSP count: 8 (L1: 4, L2 4, local L1: 2, local L2 2)

```

This table describes the significant fields shown in the display.

Table 62: show isis checkpoint lsp Field Descriptions

| Field | Description |
|----------|---|
| Level | Routers with Level 1 or Level 2 adjacency configured. |
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is 0 , the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.</p> |
| Chkpt ID | Unique ID issued to the checkpoint at its creation. |

Related Commands

| Command | Description |
|---|---|
| show isis checkpoint adjacency, on page 841 | Displays the IS-IS checkpoint adjacency database. |
| show isis database, on page 847 | Displays the IS-IS link-state database. |

show isis database

To display the Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) database, use the **show isis database** command in EXEC mode.

```
show isis [instance instance-id] database [level {1| 2}] [update] [summary] [detail] [verbose] [*| lsp-id]
```

Syntax Description

| | |
|------------------------------------|--|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS LSP database for the specified instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| level { 1 2 } | (Optional) Displays the IS-IS LSP database for Level 1 or Level 2 independently. |
| update | (Optional) Displays contents of LSP database managed by update thread. |
| summary | (Optional) Displays the LSP ID number, sequence number, checksum, hold time, and bit information. |
| detail | (Optional) Displays the contents of each LSP. |
| verbose | (Optional) Displays the contents of each LSP. |
| * <i>lsp-id</i> | (Optional) LSP protocol data units (PDUs) identifier. Displays the contents of a single LSP by its ID number or may contain an * as a wildcard character. |

Command Default

No instance ID specified displays the IS-IS LSP database for all the IS-IS instances. Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Each of the options for the **show isis database** command can be entered in an arbitrary string within the same command entry. For example, the following are both valid command specifications and provide the same output: **show isis database detail level 2** and **show isis database level 2 detail**.

The **summary** keyword used with this command allows you to filter through a large IS-IS database and quickly identify problematic areas.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis database** command with no keywords specified:

```
RP/0/RSP0/CPU0:router# show isis database

IS-IS Area a1 (Level-1) Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
router-gsr6.00-00 * 0x00000016  0x62c8        896           0/0/0
router-gsr6.01-00 * 0x0000000f  0x56d9        902           0/0/0
router-gsr8.00-00  0x00000019  0x4b6d        1015          0/0/0
router-gsr9.00-00  0x00000016  0x33b7        957           0/0/0

Total LSP count: 4 (L1: 4, L2 0, local L1: 2, local L2 0)
```

This table describes the significant fields shown in the display.

Table 63: show isis database Field Descriptions

| Field | Description |
|-------|--|
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.</p> |

| Field | Description |
|--------------|---|
| LSP Seq Num | Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source. |
| LSP Checksum | Checksum of the entire LSP packet. |
| LSP Holdtime | Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed. |
| ATT/P/OL | <p>ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router.</p> <p>P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair.</p> <p>OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.</p> |

The following is sample output from the **show isis database** command with the **summary** keyword:

```
RP/0/RSP0/CPU0:router# show isis database summary

IS-IS 10 Database Summary for all LSPs
                        Active          Purged          All
                        L1    L2    Total    L1    L2    Total    L1    L2    Total
-----
Fragment 0 Counts
  Router LSPs:         1     1     2       0     0     0       1     1     2
  Pseudo-node LSPs:   0     0     0       0     0     0       0     0     0
  All LSPs:            1     1     2       0     0     0       1     1     2
Per Topology
  IPv4 Unicast
    ATT bit set LSPs:  0     0     0       0     0     0       0     0     0
    OVL bit set LSPs:  0     0     0       0     0     0       0     0     0
All Fragment Counts
  Router LSPs:         1     1     2       0     0     0       1     1     2
  Pseudo-node LSPs:   0     0     0       0     0     0       0     0     0
  All LSPs:            1     1     2       0     0     0       1     1     2
```

This table describes the significant fields shown in the display.

Table 64: show isis database summary Field Descriptions

| Field | Description |
|-------------------|---|
| Router LSPs | Active, purged, and total LSPs associated with routers. |
| Pseudo-node LSPs: | Active, purged, and total LSPs associated with pseudonodes. |
| All LSPs: | Total active and purged LSPs. |
| ATT bit set LSPs | Attach bit (ATT). Indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router. |
| OVL bit set LSPs | Overload bit. Indicates if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router. |

The following is sample output from the **show isis database** command with the **detail** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp database detail

IS-IS isp (Level-1) Link State Database
  LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
  router-5.00-00      0x00000003  0x8074460 457              0/0/0
  Area Address: 49
  Area Address: 01
  NLPID:              0xcc
  Hostname:           router-5
  IP Address:         172.1.1.5
  Metric: 0 IP 172.3.55.0/24
  Metric: 10 IP 172.6.1.0/24
  MT:                 IPv6 Unicast
  Metric: 10 IP 172.7.0.0/24
  Metric: 10 IS router-11.00
  Metric: 10 IS router-11.01
  router-11.00-00 * 0x0000000b 0x8074460 1161              0/0/0
  Area Address: 49
  NLPID:              0xcc
  Hostname:           router-11
  IP Address:         192.168.0.145
  IP Address:         172.1.11.11 Metric: 0 IP 172.1.111.0/24
  Metric: 10 IP 172.016.1.0/24
  Metric: 10 IP 172.007.0.0/24
  Metric: 10 IS router-11.01
  Metric: 10 IS router-5.00
  router-11.01-00 * 0x00000001 0x80770ec 457              0/0/0
  Metric: 0 IS router-11.00
  Metric: 0 IS router-5.00
  Affinity: 0x00000000
  Interface IP Address: 10.3.11.145
  Neighbor IP Address: 10.3.11.143
  Physical BW: 155520 kbits/sec
```

```

Total LSP count: 3 (L1: 3, L2 0, local L1: 2, local L2 0)
Reservable Global pool BW: 0 kbits/sec
Global Pool BW Unreserved:
IS-IS isp (Level-2) Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
router-5.00-00  0x00000005  0x807997c     457           0/0/0
  [0]: 0          kbits/sec      [1]: 0          kbits/sec
  [2]: 0          kbits/sec      [3]: 0          kbits/sec
  [4]: 0          kbits/sec      [5]: 0          kbits/sec
  [6]: 0          kbits/sec      [7]: 0          kbits/sec
Area Address: 49
Interface IP Address: 10.3.11.145
Neighbor IP Address: 10.3.11.143
  NLPID: 0xcc    Hostname: router-5IP Address: 172.6.1.5
  Metric: 0     IP 172.3.55.0/24
  Metric: 10    IP 172.1686.1.0/24
  Metric: 10    IS router-11.00
  Metric: 10    IP 172.1.0.0/24
  Metric: 10    IS router-11.01
  Metric: 10    IP 172.8.111.0/24

router-11.00-00 * 0x0000000d  0x807997c     1184           0/0/0
Area Address: 49
NLPID: 0xcc
Hostname: router-11
IP Address: 172.28.111.111
Metric: 0     IP 172.8.111.0/24
Metric: 10    IP 172.6.1.0/24
Metric: 10    IP 172.7.0.0/24
Metric: 10    IS router-11.01
Metric: 10    IS router-5.00
Metric: 10    IP 172.3.55.0/24  router-gsr11.01-00 * 0x00000001 0x80770ec
457          0/0/0
Metric: 0     IS router-11.00
Metric: 0     IS router-5.00

Total LSP count: 3 (L1: 0, L2 3, local L1: 0, local L2 2)

```

As the output shows, besides the information displayed with the **show isis database** command, the command with the **detail** keyword displays the contents of each LSP.

Table 65: show isis instance isp database detail Field Descriptions

| Field | Description |
|--------------|--|
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.</p> |
| LSP Seq Num | Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source. |
| LSP Checksum | Checksum of the entire LSP packet. |
| LSP Holdtime | Amount of time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed. |

| Field | Description |
|--------------|---|
| ATT/P/OL | <p>ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router.</p> <p>P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair.</p> <p>OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router.</p> |
| Area Address | Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area this route belongs to. |
| NLPID | Network Layer Protocol Identifier. |
| Hostname | Hostname of the node. |
| IP Address: | Address of the node. |
| Metric | IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a Connectionless Network Service (CLNS) prefix). |

The following is additional sample output from the **show isis database detail** command. This is a Level 2 LSP. The area address 39.0001 is the address of the area in which the router resides.

```
RP/0/RSP0/CPU0:router# show isis database level 2 detail

IS-IS Level-2 Link State Database
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0000.0C00.1111.00-00* 0x00000006   0x4DB3        1194          0/0/0
Area Address: 39.0001
NLPID:        0x81 0xCC
IP Address:   172.18.1.17
Metric: 10    IS 0000.0C00.1111.09
Metric: 10    IS 0000.0C00.1111.08
Metric: 10    IP 172.17.4.0 255.255.255.0
Metric: 10    IP 172.18.8.0 255.255.255.0
Metric: 0     IP-External 10.0.0.0 255.0.0.0
```

The IP entries are the directly connected IP subnets the router is advertising (with associated metrics). The IP-External entry is a redistributed route.

Table 66: show isis database level 2 detail Field Descriptions

| Field | Description |
|--------------|--|
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.</p> |
| LSP Seq Num | Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source. |
| LSP Checksum | Checksum of the entire LSP packet. |
| LSP Holdtime | Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed. |

| Field | Description |
|--------------|--|
| ATT/P/OL | <p>ATT—Attach bit. This bit indicates that the router is also a Level 2 router, and it can reach other areas. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers use the Attach bit to find the closest Level 2 router. They point to a default route to the closest Level 2 router.</p> <p>P—P bit. Detects if the intermediate system is area partition repair capable. Cisco and other vendors do not support area partition repair.</p> <p>OL—Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routes. Only packets for destinations directly connected to the overloaded router are sent to this router.</p> |
| Area Address | Reachable area addresses from the router. For Level 1 LSPs, these are the area addresses configured manually on the originating router. For Level 2 LSPs, these are all the area addresses for the area to which this route belongs. |
| NLPID | Network Layer Protocol Identifier. |
| Hostname | Hostname of the node. |
| IP Address: | IP address of the node. |
| Metric: | IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a Connectionless Network Service (CLNS) prefix). |

show isis database-log

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) database log, use the **show isis database-log** command in EXEC mode.

show isis database-log [**level** {**1** | **2**}] [**last number** | **first number**]

Syntax Description

| | |
|--------------------------------------|---|
| level { 1 2 } | (Optional) Displays the database log for Level 1 or Level 2 independently. |
| last number | (Optional) Specifies that the output be restricted to the last <i>number</i> of entries. Range is 1 to 1000. |
| first number | (Optional) Specifies that the output be restricted to the first <i>number</i> of entries. Range is 1 to 1000. |

Command Default

Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis database-log** command:

```
RP/0/RSP0/CPU0:router# show isis database-log

IS-IS 10 Level 1 Link State Database Log
                New LSP
WHEN      LSPID                Op  Seq Num  Holdtime OL  Old LSP
01:17:19 12b1.03-00             REP 0x00000003 1200   0 0x00000002 340   0
001:06:20 12b1.00-00             REP 0x000001d8 1200   0 0x000001d7 375   0
```

```

01:06:00 12b1.03-00      REP 0x00000004 1200      0 0x00000003 520      0
01:05:46 12a1.00-00      REP 0x000001fc 1200      0 0x000001fb 425      0
00:55:01 12b1.00-00      REP 0x000001d9 1200      0 0x000001d8 520      0
00:53:39 12b1.03-00      REP 0x00000005 1200      0 0x00000004 459      0
00:53:19 12a1.00-00      REP 0x000001fd 1200      0 0x000001fc 453      0
00:42:12 12b1.00-00      REP 0x000001da 1200      0 0x000001d9 431      0
00:39:56 12b1.03-00      REP 0x00000006 1200      0 0x00000005 376      0
00:38:54 12a1.00-00      REP 0x000001fe 1200      0 0x000001fd 334      0
00:29:10 12b1.00-00      REP 0x000001db 1200      0 0x000001da 418      0
00:27:22 12b1.03-00      REP 0x00000007 1200      0 0x00000006 446      0
00:25:10 12a1.00-00      REP 0x000001ff 1200      0 0x000001fe 375      0
00:17:04 12b1.00-00      REP 0x000001dc 1200      0 0x000001db 473

```

This table describes the significant fields shown in the display.

Table 67: show isis database-log Field Descriptions

| Field | Description |
|---------|--|
| WHEN | Elapsed time (in hh:mm:ss) since the event was logged. |
| LSPID | <p>LSP identifier. The first six octets form the system ID of the router that originated the LSP.</p> <p>The next octet is the pseudonode ID. When this byte is 0, the LSP describes links from the system. When it is nonzero, the LSP is a so-called nonpseudonode LSP. This is similar to a router link-state advertisement (LSA) in the Open Shortest Path First (OSPF) protocol. The LSP describes the state of the originating router.</p> <p>For each LAN, the designated router for that LAN creates and floods a pseudonode LSP, describing all systems attached to that LAN.</p> <p>The last octet is the LSP number. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments. Each fragment has a different LSP number. An asterisk (*) indicates that the LSP was originated by the system on which this command is issued.</p> |
| New LSP | New router or pseudonode appearing in the topology. |
| Old LSP | Old router or pseudonode leaving the topology. |
| Op | Operation on the database: inserted (INS) or replaced (REP). |
| Seq Num | Sequence number for the LSP that allows other systems to determine if they have received the latest information from the source. |

| Field | Description |
|----------|---|
| Holdtime | Time the LSP remains valid (in seconds). An LSP hold time of 0 indicates that this LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP stays in the LSDB before being completely removed. |
| OL | Overload bit. Determines if the IS is congested. If the Overload bit is set, other routers do not use this system as a transit router when calculating routers. Only packets for destinations directly connected to the overloaded router are sent to this router. |

Related Commands

| Command | Description |
|---|--|
| show isis database, on page 847 | Displays the IS-IS link-state packet (LSP) database. |

show isis fast-reroute

To display per-prefix LFA information, use the **show isis fast-reroute** command in EXEC mode.

show isis fast-reroute *A.B.C.D/length* | **detail** | **summary**

| Syntax Description | |
|-----------------------|---|
| <i>A.B.C.D/length</i> | Network to show per-prefix LFA information. |
| detail | Use to display tiebreaker information about the backup. |
| summary | Use to display the number of prefixes having protection per priority. |

Command Default None

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.0.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | isis | read |

Examples The following is sample output from **show isis fast-reroute** command that displays per-prefix LFA information:

```
RP/0/RSP0/CPU0:router# show isis fast-reroute 10.1.6.0/24
L1 10.1.6.0/24 [20/115]
   via 10.3.7.47, POS0/3/0/1, router2
   FRR backup via 10.1.7.145, GigabitEthernet0/1/0/3, router3
```

The following is sample output from **show isis fast-reroute detail** command that displays tie-breaker information about the backup:

```
RP/0/RSP0/CPU0:router# show isis fast-reroute 10.1.6.0/24 detail
L1 10.1.6.0/24 [20/115] low priority
   via 10.3.7.47, POS0/3/0/1, router2
   FRR backup via 10.1.7.145, GigabitEthernet0/1/0/3, router3
   P: No, TM: 30, LC: Yes, NP: No, D: No
```

```
src router2.00-00, 192.168.0.47
L2 adv [20] native, propagated
```

The following is sample output from **show isis fast-reroute summary** command that displays the number of prefixes having protection per priority:

```
RP/0/RSP0/CPU0:router#show isis fast-reroute summary
IS-IS frf IPv4 Unicast FRR summary
```

| | | Critical Priority | High Priority | Medium Priority | Low Priority | Total |
|--------------------------|-------|----------------------|------------------|--------------------|-----------------|-------|
| Prefixes reachable in L1 | | | | | | |
| All paths protected | 0 | 0 | 2 | | 8 | 10 |
| Some paths protected | 0 | 0 | 1 | | 3 | 4 |
| Unprotected | 0 | 0 | 1 | | 3 | |
| Protection coverage | 0.00% | 0.00% | 75.00% | 78.57% | 77.78% | |
| Prefixes reachable in L2 | | | | | | |
| All paths protected | 0 | 0 | 0 | | 0 | 0 |
| Some paths protected | 0 | 0 | 1 | | 0 | 1 |
| Unprotected | 0 | 0 | 0 | | 0 | |
| Protection coverage | 0.00% | 0.00% | 100.00% | 0.00% | 100.00% | |

show isis hostname

To display the entries in the Intermediate System-to-Intermediate System (IS-IS) router name-to-system ID mapping table, use the **show isis hostname** command in EXEC mode.

show isis [*instance instance-id*] **hostname**

Syntax Description

instance *instance-id* (Optional) Displays the IS-IS router name-to-system ID mapping table for the specified IS-IS instance only.

The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS router name-to-system ID mapping table for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show isis hostname** command does not display entries if the dynamic hostnames are disabled.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis hostname** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp hostname
ISIS isp hostnames
  Level System ID      Dynamic Hostname
  1      0001.0000.0005 router
  2      * 0001.0000.0011 router-11
```

This table describes the significant fields shown in the display.

Table 68: show isis instance isp hostname Field Descriptions

| Field | Description |
|------------------|--|
| Level | IS-IS level of the router. |
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| Dynamic Hostname | Hostname of the router. |
| * | Local router. |

Related Commands

| Command | Description |
|---|---|
| hostname | Specifies the name of the local router. |
| hostname dynamic disable, on page 749 | Enables the IS-IS routing protocol to dynamically update the mapping of router names to system IDs. |

show isis interface

To display information about the Intermediate System-to-Intermediate System (IS-IS) interfaces, use the **show isis interface** command in EXEC mode.

show isis interface [*type interface-path-id*] **level** {1| 2} [**brief**]

Syntax Description

| | |
|-------------------|---|
| type | Interface type. For more information, use the question mark (?) online help function. |
| interface-path-id | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| level { 1 2 } | (Optional) Displays IS-IS interface information for Level 1 or Level 2 independently. |
| brief | (Optional) Displays brief interface output. |

Command Default

Displays all IS-IS interfaces.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis interface** command:

```
RP/0/RSP0/CPU0:router#show isis interface
      GigabitEthernet 0/3/0/2
Gi /3/0/2      Enabled
Adjacency Formation:      Enabled
Prefix Advertisement:      Enabled
BFD:                  Disabled
BFD Min Interval:        150
BFD Multiplier:          3

Circuit Type:            level-2-only
Media Type:              P2P
Circuit Number:          0
Extended Circuit Number: 67111168
Next P2P IIH in:         4 s
LSP Rermit Queue Size:   0

Level-2
Adjacency Count:         1
LSP Pacing Interval:     33 ms
PSNP Entry Queue Size:   0

CLNS I/O
Protocol State:          Up
MTU:                     4469

IPv4 Unicast Topology:   Enabled
Adjacency Formation:     Running
Prefix Advertisement:    Running
Metric (L1/L2):          10/100
MPLS LDP Sync (L1/L2):  Disabled/Disabled
IPv6 Unicast Topology:   Disabled (Not cfg on the intf)

IPv4 Address Family:     Enabled
Protocol State:           Up
Forwarding Address(es):  10.3.10.143
Global Prefix(es):       10.3.10.0/24
IPv6 Address Family:     Disabled (No topology enabled which uses IPv6)

LSP transmit timer expires in 0 ms
LSP transmission is idle
Can send up to 9 back-to-back LSPs in the next 0 ms
```

This table describes the significant fields shown in the display.

Table 69: show isis interface Field Descriptions

| Field | Description |
|------------------------|---|
| GigabitEthernet0/6/0/0 | Status of the interface, either enabled or disabled. |
| Adjacency formation | Status of adjacency formation, either enabled or disabled. |
| Prefix Advertisement | Status of advertising connected prefixes, either enabled or disabled. |
| BFD | Status of Bidirectional Forwarding Detection (BFD), either enabled or disabled. |

| Field | Description |
|-------------------------|--|
| BFD Min Interval | BFD minimum interval. |
| BFD Multiplier | BFD multiplier. |
| Circuit Type | Levels the interface is running on (circuit-type configuration) which may be a subset of levels on the router. |
| Media Type | Media type on which IS-IS is running. |
| Circuit Number | Unique ID assigned to a circuit internally (8-bit integer). |
| Extended Circuit Number | Valid only for point-to-point interfaces (32-bit integer). |
| LSP Reremit Queue Size | Number of LSPs pending retransmission on the interface. |
| Adjacency Count | Number of adjacencies formed with a neighboring router that supports the same set of protocols. |
| PSNP Entry Queue Size | Number of SNP entries pending inclusion in the next PSNP. |
| LAN ID | ID of the LAN. |
| Priority (Local/DIS) | Priority of this interface or priority of the Designated Intermediate System. |
| Next LAN IIH in | Time (in seconds) in which the next LAN hello message is sent. |
| LSP Pacing Interval | Interval at which the link-state packet (LSP) transmission rate (and by implication the reception rate of other systems) is to be reduced. |
| Protocol State | Running state of the protocol (up or down). |
| MTU | Link maximum transmission unit (MTU). |
| SNPA | Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor. |
| All Level-n ISs | Status of interface membership in Layer 2 multicast group. The status options are Yes or reason for not being a member of the multicast group. |
| IPv4 Unicast Topology | Status of the topology, either enabled or disabled. |

| Field | Description |
|-------------------------------|--|
| Adjacency Formation | Status of adjacency formation. The status options are Running or a reason for not being ready to form adjacencies. |
| Prefix Advertisement | Status of advertising prefixes, either enabled or disabled. |
| Metric (L1/L2) | IS-IS metric for the cost of the adjacency between the originating router and the advertised neighbor, or the metric of the cost to get from the advertising router to the advertised destination (which can be an IP address, an end system (ES), or a connectionless network service (CLNS) prefix). |
| MPLS LDP Sync (L1/L2) | Status of LDP IS-IS synchronization, either enabled or disabled. When enabled, the state of synchronization (Sync Status) is additionally displayed as either achieved or not achieved. |
| IPv4 Address Family | Status of the address family, either enabled or disabled. |
| Protocol State | State of the protocol. |
| Forwarding Address(es) | Addresses on this interface used by the neighbor for next-hop forwarding. |
| Global Prefix(es) | Prefixes for this interface included in the LSP. |
| LSP transmit timer expires in | LSP transmission expiration timer interval (in milliseconds). |
| LSP transmission is | State of LSP transmission. Valid states are <ul style="list-style-type: none"> • idle • in progress • requested • requested and in progress |

The following is sample output from the **show isis interface** command with the **brief** keyword:

```
RP/0/0/CPU0:router# show isis interface brief
```

| Interface | All | Adjs | | Adj Topos | Adv Topos | CLNS | MTU | Prio | |
|-----------|-----|------|----|-----------|-----------|------|------|------|----|
| | OK | L1 | L2 | Run/Cfg | Run/Cfg | | | L1 | L2 |
| PO0/5/0/0 | Yes | 1 | 1 | 1/1 | 1/1 | Up | 4469 | - | - |
| Gi0/6/0/0 | Yes | 1* | 1* | 1/1 | 1/1 | Up | 1497 | 64 | 64 |

This table describes the significant fields shown in the display.

Table 70: show isis interface brief Field Descriptions

| Field | Description |
|-------------------|--|
| Interface | Name of the interface. |
| All OK | Everything is working as expected for this interface. |
| Adjs L1 L2 | Number of L1 and L2 adjacencies over this interface. |
| Adj Topos Run/Cfg | Number of topologies that participate in forming adjacencies. Number of topologies that were configured to participate in forming adjacencies. |
| Adv Topos Run/Cfg | Number of topologies that participate in advertising prefixes. Number of topologies that were configured to participate in advertising prefixes. |
| CLNS | Status of the Connectionless Network Service. Status options are Up or Down. |
| MTU | Maximum transfer unit size for the interface. |
| Prio L1 L2 | Interface L1 priority. Interface L2 priority. |

show isis lsp-log

To display link-state packet (LSP) log information, use the **show isis lsp-log** command in EXEC mode.

```
show isis [instance instance-id] lsp-log [level {1| 2}] [last number| first number]
```

Syntax Description

| | |
|------------------------------------|--|
| instance <i>instance-id</i> | (Optional) Displays the LSP log information for the specified IS-IS instance only. • The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| level { 1 2 } | (Optional) Displays the Intermediate System-to-Intermediate System (IS-IS) link-state database for Level 1 or Level 2 independently. |
| last <i>number</i> | (Optional) Specifies that the output be restricted to the last <i>number</i> of entries. Range is 1 to 20. |
| first <i>number</i> | (Optional) Specifies that the output be restricted to the first <i>number</i> of entries. Range is 1 to 20. |

Command Default

No instance ID specified displays the LSP log information for all the IS-IS instances. Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis lsp-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp lsp-log

ISIS isp Level 1 LSP log
  When          Count      Interface      Triggers
00:02:36         1
00:02:31         1
00:02:26         1      PO4/1         LSPREGEN
00:02:24         1      PO4/1         DELADJ
00:02:23         1      PO4/1         NEWADJ
00:02:23         1      Gi5/0         DIS
00:01:27         1      Lo0           IPDOWN
00:01:12         1      Lo0           IPUP

ISIS isp Level 2 LSP log
  When          Count      Interface      Triggers
00:02:36         1
00:02:30         1
00:02:26         1      PO4/1         LSPREGEN
00:02:24         1      PO4/1         DELADJ
00:02:23         1      PO4/1         NEWADJ
00:02:21         1      Gi5/0         DIS
00:02:21         1      Gi5/0         AREASET
00:01:27         1      Lo0           IPDOWN
00:01:12         1      Lo0           IPUP
```

This table describes the significant fields shown in the display.

Table 71: show isis instance isp lsp-log Field Descriptions

| Field | Description |
|-----------|--|
| Level | IS-IS level of the router. |
| When | How long ago (in hh:mm:ss) an LSP rebuild occurred. The last 20 occurrences are logged. |
| Count | Number of events that triggered this LSP run. When there is a topology change, often multiple LSPs are received in a short period. A router waits 5 seconds before running a full LSP, so it can include all new information. This count denotes the number of events (such as receiving new LSPs) that occurred while the router was waiting its 5 seconds before running full LSP. |
| Interface | Interface that corresponds to the triggered reasons for the LSP rebuild. |

| Field | Description |
|----------|---|
| Triggers | <p>A list of all reasons that triggered an LSP rebuild. The triggers are</p> <ul style="list-style-type: none"> • AREASET—area set changed • ATTACHFLAG—bit attached • CLEAR— clear command • CONFIG—configuration change • DELADJ—adjacency deleted • DIS—DIS changed • IFDOWN—interface down • IPADDRCHG—IP address change • IPDEFORIG—IP def-orig • IPDOWN—connected IP down • IFDOWN—interface down • IPEXT—external IP • IPIA—nterarea IP • IPUP—connected IP up • LSPDBOL—LSPDBOL bit • LSPREGEN—LSP regeneration • NEWADJ— new adjacency |

show isis mesh-group

To display Intermediate System-to-Intermediate System (IS-IS) mesh group information, use the **show isis mesh-group** command in EXEC mode.

show isis [*instance instance-id*] **mesh-group**

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Displays the mesh group information for the specified IS-IS instance only. |
|------------------------------------|---|

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the IS-IS mesh group information for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis mesh-group** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mesh-group
ISIS isp Mesh Groups
Mesh group 6:
GigabitEthernet 0/4/0/1
```

This table describes the significant fields shown in the display.

Table 72: show isis instance isp mesh-group Field Descriptions

| Field | Description |
|------------------------|---|
| Mesh group | Mesh group number to which this interface is a member. A mesh group optimizes link-state packet (LSP) flooding in nonbroadcast multiaccess (NBMA) networks with highly meshed, point-to-point topologies. LSPs that are first received on interfaces that are part of a mesh group are flooded to all interfaces except those in the same mesh group. |
| GigabitEthernet0/4/0/1 | Interface belonging to mesh group 6. |

show isis mpls traffic-eng adjacency-log

To display a log of Multiprotocol Label Switching traffic engineering (MPLS TE) adjacency changes for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng adjacency-log** command in EXEC mode.

```
show isis [instance instance-id] mpls traffic-eng adjacency-log [last number| first number]
```

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Displays the MPLS TE adjacency changes for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| last <i>number</i> | (Optional) Specifies that the output is restricted to last <i>number</i> of entries. Range is 1 to 20. |
| first <i>number</i> | (Optional) Specifies that the output is restricted to first <i>number</i> of entries. Range is 1 to 20. |

Command Default

No instance ID specified displays MPLS TE adjacency changes for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis mpls traffic-eng adjacency-log** command to display the status of MPLS TE adjacencies.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis mpls traffic-eng adjacency-log** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mpls traffic-eng adjacency-log

IS-IS isp Level-2 MPLS Traffic Engineering adjacency log
When      Neighbor ID      IP Address      Interface Status
00:03:36  router-6        172.17.1.6     PO0/3/0/1 Up
00:03:36  router-6        172.17.1.6     PO0/3/0/1 Down
00:02:38  router-6        172.17.1.6     PO0/3/0/1 Up
```

This table describes the significant fields shown in the display.

Table 73: show isis instance isp mpls traffic-eng adjacency-log Field Descriptions

| Field | Description |
|-------------|---|
| When | Time (in hh:mm:ss) since the entry was recorded in the log. |
| Neighbor ID | Identification value of the neighbor. |
| IP Address | Neighbor IP Version 4 (IPv4) address. |
| Interface | Interface from which a neighbor is learned. |
| Status | Up (active) or Down (disconnected). |

Related Commands

| Command | Description |
|---|---|
| show isis mpls traffic-eng advertisements , on page 875 | Displays the last flooded record from MPLS traffic engineering. |

show isis mpls traffic-eng advertisements

To display the latest flooded record from Multiprotocol Label Switching traffic engineering (MPLS TE) for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng advertisements** command in EXEC mode.

show isis [*instance instance-id*] **mpls traffic-eng advertisements**

Syntax Description

instance *instance-id* (Optional) Displays the latest flooded record from MPLS TE for the specified IS-IS instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the latest flooded record from MPLS TE for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis mpls traffic-eng advertisements** command to verify that MPLS TE is flooding its record and that the bandwidths are correct.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis mpls traffic-eng advertisements** command with the **instance** and *instance-id* values specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp mpls traffic-eng advertisements
ISIS isp Level-2 MPLS Traffic Engineering advertisements
```

show isis mpls traffic-eng advertisements

```

System ID: router-9
Router ID: 172.18.0.9
Link Count: 1
  Link[0]
    Neighbor System ID: router-gsr6 (P2P link)
    Interface IP address: 172.18.0.9
    Neighbor IP Address: 172.18.0.6
    Admin. Weight: 0
    Physical BW: 155520000 bits/sec
    Reservable BW global: 10000000 bits/sec
    Reservable BW sub: 0 bits/sec
    Global pool BW unreserved:
      [0]: 10000000 bits/sec, [1]: 10000000 bits/sec
      [2]: 10000000 bits/sec, [3]: 10000000 bits/sec
      [4]: 10000000 bits/sec, [5]: 10000000 bits/sec
      [6]: 10000000 bits/sec, [7]: 10000000 bits/sec
    Sub pool BW unreserved:
      [0]: 0 bits/sec, [1]: 0 bits/sec
      [2]: 0 bits/sec, [3]: 0 bits/sec
      [4]: 0 bits/sec, [5]: 0 bits/sec
      [6]: 0 bits/sec, [7]: 0 bits/sec
    Affinity Bits: 0x00000000

```

This table describes the significant fields shown in the display.

Table 74: show isis instance isp mpls traffic-eng advertisements Field Descriptions

| Field | Description |
|----------------------|--|
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or if the hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| Router ID | MPLS TE router ID. |
| Link Count | Number of links that MPLS TE advertised. |
| Neighbor System ID | System ID of a neighbor number in an area. The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2). In an IS-IS routing domain, each router is represented by a 6-byte hexadecimal system ID. When network administrators maintain and troubleshoot networking devices, they must know the router name and corresponding system ID. |
| Interface IP address | IP address of the interface. |
| Neighbor IP Address | IP address of the neighbor. |
| Admin. Weight | Administrative weight associated with this link. |
| Physical BW | Link bandwidth capacity (in bits per second). |

| Field | Description |
|---------------------------|--|
| Reservable BW | Reservable bandwidth on this link. |
| Global pool BW unreserved | Unreserved bandwidth that is available in the global pool. |
| Sub pool BW unreserved | Amount of unreserved bandwidth that is available in the subpool. |
| Affinity Bits | Link attribute flags being flooded. Bits are MPLS-TE specific. |

Related Commands

| Command | Description |
|--|--|
| show isis mpls traffic-eng adjacency-log , on page 873 | Displays a log of MPLS TE adjacency changes for IS-IS. |

show isis mpls traffic-eng tunnel

To display Multiprotocol Label Switching traffic engineering (MPLS TE) tunnel information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis mpls traffic-eng tunnel** command in EXEC mode.

show isis [*instance instance-id*] **mpls traffic-eng tunnel**

Syntax Description

instance *instance-id* (Optional) Displays the MPLS TE tunnel information for the specified IS-IS instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays the MPLS TE tunnel information for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis** command to find the current status of MPLS TE tunnels.

Tunnels are used in IS-IS next-hop calculations.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis mpls traffic-eng tunnel** command:

```
RP/0/RSP0/CPU0:router# show isis mpls traffic-eng tunnel
ISIS isp Level-2 MPLS Traffic Engineering tunnels
System Id          Tunnel Name      Bandwidth  Nexthop      Metric  Mode
```

```
router-6          tu0          100000          172.18.1.6          0          Relative
```

This table describes the significant fields shown in the display.

Table 75: show isis mpls traffic-eng tunnel Field Descriptions

| Field | Description |
|-------------|--|
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| Tunnel Name | Name of the MPLS TE tunnel interface. |
| Bandwidth | MPLS TE-specified tunnel bandwidth of the tunnel. |
| Nexthop | MPLS TE destination IP address of the tunnel. |
| Metric | MPLS TE metric of the tunnel. |
| Mode | MPLS TE metric mode of the tunnel. It can be relative or absolute. |

show isis neighbors

To display information about Intermediate System-to-Intermediate System (IS-IS) neighbors, use the **show isis neighbors** command in EXEC mode.

show isis [**instance** *instance-id*] **neighbors** [*type interface-path-id*] **summary**] [**detail**] [**systemid** *system-id*]

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS neighbor information for the specified IS-IS instance only. • The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| type | Interface type. For more information, use the question mark (?) online help function. |
| interface-path-id | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| summary | (Optional) Displays neighbor status count for each level. |
| detail | (Optional) Displays additional details. |
| systemid <i>system-id</i> | (Optional) Displays the information for the specified neighbor only. |

Command Default

No instance ID specified displays neighbor information for all the IS-IS instances.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis neighbors** command with the **instance** and *instance-id* values specified:

```
Total neighbor count: 3
RP/0/RSP0/CPU0:router# show isis instance isp neighbors detail

IS-IS isp neighbors:
System Id      Interface      SNPA          State Holdtime Type IETF-NSF
e222e          Gi0/1/0/0     *PtoP*       Up    23      L1    Capable
  Area Address(es): 00
  IPv4 Address(es): 10.1.0.45*
  IPv6 Address(es): fe80::212:daff:fe6b:68a8*
  Topologies: 'IPv4 Unicast'
  Uptime: 01:09:44
  IPFRR: LFA Neighbor: elise
          LFA IPv4 address: 10.100.1.2
          LFA Router address: 192.168.0.45
e333e Gi0/1/0/0.1  0012.da6b.68a8 Up    8      L1    Capable
e333e          Gi0/1/0/0.1  0012.da6b.68a8 Up    8      L1    Capable
  Area Address(es): 00
  IPv4 Address(es): 10.100.1.2*
  Topologies: 'IPv4 Unicast'
  Uptime: 01:09:46
  IPFRR: LFA Neighbor: elise
          LFA IPv4 address: 10.1.0.45
          LFA Router address: 192.168.0.45
          LFA Interface: Gi0/1/0/0
m44i          Gi0/1/0/1     0012.da62.e0a8 Up    7      L1    Capable
  Area Address(es): 00 11
  IPv4 Address(es): 10.1.2.47*
  IPv6 Address(es): fe80::212:daff:fe62:e0a8*
  Topologies: 'IPv4 Unicast'
  Uptime: 01:09:33

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 76: show isis instance isp neighbors Field Descriptions

| Field | Description |
|-----------|--|
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| Interface | Interface through which the neighbor is reachable. |
| SNPA | Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor. |

| Field | Description |
|---------------------|---|
| State | Adjacency state of the neighboring interface. Valid states are: Down, Init, and Up. |
| Holdtime | Hold time of the neighbor. |
| Type | Type of adjacency. |
| IETF-NSF | Specifies whether the neighbor can adhere to the IETF-NSF restart mechanism. Valid states are Capable and Unable. |
| Area Address(es) | Number of area addresses on this router. |
| IPv4 Address(es) | IPv4 addresses configured on this router. |
| Topologies | Address and subaddress families for which IS-IS is configured. |
| Uptime | Time (in hh:mm:ss) that the neighbor has been up. |
| IPFRR: LFA Neighbor | IP fast reroute (IPFRR) loop-free alternate (LFA) neighbor. |
| LFA IPv4 address: | Address of the LFA. |
| LFA Interface: | LFA interface. |

The following is sample output from the **show isis neighbors** command with the **summary** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis instance isp neighbors summary

ISIS isp neighbor summary:
  State      L1    L2    L1L2
  Up         0     0     2
  Init       0     0     0
  Failed     0     0     0
```

This table describes the significant fields shown in the display.

Table 77: show isis neighbors summary Field Descriptions

| Field | Description |
|-------|--|
| State | State of the neighbor is up, initialized, or failed. |
| L1 | Number of Level 1 neighbors. |
| L2 | Number of Level 2 neighbors. |

| Field | Description |
|-------|------------------------------------|
| L1L2 | Number of Level 1 and 2 neighbors. |

Related Commands

| Command | Description |
|--|-----------------------------|
| show isis adjacency, on page 836 | Displays IS-IS adjacencies. |

show isis protocol

To display summary information about an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis protocol** command in EXEC mode.

show isis [**instance** *instance-id*] **protocol**

Syntax Description

instance *instance-id* (Optional) Displays the IS-IS adjacencies for the specified IS-IS instance only.

- The *instance-id* argument is the instance identifier (alphanumeric) defined by the **router isis** command.

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis protocol** command:

```
RP/0/RSP0/CPU0:router# show isis protocol

IS-IS Router: isp
System Id: 0001.0000.0011
IS Levels: level-1-2
Manual area address(es):
  49

Routing for area address(es):
  49
Non-stop forwarding: Cisco Proprietary NSF Restart enabled
```

```

Process startup mode: Cold Restart
Topologies supported by IS-IS:
  IPv4 Unicast
    Level-1 iSPF status: Dormant (awaiting initial convergence)
    Level-2 iSPF status: Dormant (awaiting initial convergence)
    No protocols redistributed
    Distance: 115
Interfaces supported by IS-IS:
  Loopback0 is running passively (passive in configuration)
  GigabitEthernet 0/4/0/1 is running actively (active in configuration)
  GigabitEthernet 0/5/0/1 is running actively (active in configuration)

```

This table describes the significant fields shown in the display.

Table 78: show isis protocol Field Descriptions

| Field | Description |
|-----------------------------|--|
| System ID: | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| IS Levels: | IS-IS level of the router. |
| Manual area address(es) | Area addresses that are manually configured. |
| Routing for areaaddress(es) | Area addresses for which this router provides the routing. |
| Non-stop forwarding: | Status and name of nonstop forwarding (NSF). |
| Process startup mode: | Mode in which the last process startup occurred. Valid modes are: <ul style="list-style-type: none"> • Cisco Proprietary NSF Restart • IETF NSF Restart • Cold Restart |

| Field | Description |
|-----------------------------|--|
| iSPF status: | <p>State of incremental shortest path first (iSPF) configuration for this IS-IS instance. Four states exist:</p> <p>Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm.</p> <p>Dormant if iSPF has been configured but is awaiting initial convergence before initializing.</p> <p>Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm.</p> <p>Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run.</p> |
| No protocols redistributed: | No redistributed protocol information exists to be displayed. |
| Distance: | Administrative distance for this protocol. |

show isis route

To display IP reachability information for an Intermediate System-to-Intermediate System (IS-IS) instance, use the **show isis route** command in EXEC mode.

```
show isis [instance instance-id] [ipv4|ipv6] afi-all] [unicast|multicast [topology {all|topo-name}]] safi-all]
route [ip-address mask|ip-address/length [longer-prefixes]] [summary] [multicast-intact] [backup] [detail]
```

Syntax Description

| | |
|------------------------------------|--|
| instance <i>instance-id</i> | (Optional) Displays the IP reachability information for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| afi-all | (Optional) Specifies all address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| topology | (Optional) Specifies IS-IS paths to intermediate systems. |
| all | (Optional) Specifies all topologies. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| safi-all | (Optional) Specifies all secondary address prefixes. |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address. |
| <i>/ length</i> | (Optional) Length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. Range is 0 to 32. |

| | |
|-------------------------|--|
| longer-prefixes | (Optional) Displays route and more-specific routes. |
| summary | (Optional) Displays topology summary information. |
| multicast-intact | (Optional) Displays multicast intact information for this entry. |
| systemid | (Optional) Displays multicast information by system ID. |
| backup | (Optional) Displays backup information for this entry. |
| detail | (Optional) Displays link-state packet (LSP) details. |

Command Default No instance ID specified displays the IP reachability information for all the IS-IS instances.

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |
| | Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | isis | read |

Examples The following is sample output from the **show isis route** command:

```
RP/0/RSP0/CPU0:router# show isis route

IS-IS isp IPv4 Unicast routes
Codes: L1 - level 1, L2 - level 2, ia - interarea (leaked into level 1)
df - level 1 default (closest attached router), su - summary null
C - connected, S - static, R - RIP, B - BGP, O - OSPF
i - IS-IS (redistributed from another instance)

Maximum parallel path count: 8

L2 10.76.240.6/32 [4/115]
via 10.76.245.252, SRP0/1/0/2, isp2
```

```

via 10.76.246.252, SRP0/1/0/0, isp2
C 10.76.240.7/32
is directly connected, Loopback0
L2 10.76.240.9/32 [256/115]
via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3
L2 10.76.240.10/32 [296/115]
via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3
C 10.76.245.0/24
is directly connected, SRP0/1/0/2
C 10.76.246.0/24
is directly connected, SRP0/1/0/0
C 10.76.249.0/26
is directly connected, GigabitEthernet 0/3/0/0
L2 10.101.10.0/24 [296/115]
via 10.76.249.2, GigabitEthernet 0/3/0/0, isp3

```

This table describes the significant fields shown in the display.

Table 79: show isis route ipv4 unicast Field Descriptions

| Field | Description |
|-----------------------|--|
| C172.18.0.0/24 | Connected route for GigabitEthernet interface 0/5/0/0. |
| C 172.19.1.0/24 | Connected route for GigabitEthernet interface 0/4/0/1. |
| L1 172.35.0.0/24 [10] | Level 1 route to network 172.35.0.0/24. |
| C 172.18.0/24 | Connected route for loopback interface 0. |

show isis spf-log

To display how often and why the router has run a full shortest path first (SPF) calculation, use the **show isis spf-log** command in EXEC mode.

```
show isis [instance instance-id] [[ipv4|ipv6] afi-all] [unicast|multicast [topology {all|topo-name}]] safi-all]
spf-log [level {1|2}] [ispf|fspf] prc nhc] [detail|verbose|plfrr|ppfrr] [last number|first number]
```

Syntax Description

| | |
|---|---|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS SPF log for the specified IS-IS instance only. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| afi-all | (Optional) Specifies all address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| topology all <i>topo-name</i> | (Optional) Specifies topology table information for all topologies or for the specified topology table (<i>topo-name</i>). |
| safi-all | (Optional) Specifies all secondary address prefixes. |
| level { 1 2 } | (Optional) Displays the IS-IS SPF log for Level 1 or Level 2 independently. |
| ispf | (Optional) Specifies incremental SPF entries only. |
| fspf | (Optional) Specifies full SPF entries only. |
| prc | (Optional) Specifies partial route calculations only. |
| nhc | (Optional) Specifies next-hop route calculations only. |
| detail | (Optional) Specifies detailed output. Includes a breakdown of the time taken to perform the calculation and changes resulting from the calculation. |
| verbose | (Optional) Specifies verbose output. |
| last <i>number</i> | (Optional) Specifies that the output is restricted to the last <i>number</i> of entries. Range is 1 to 210. |
| first <i>number</i> | (Optional) Specifies that the output is restricted to the first <i>number</i> of entries. Range is 1 to 210. |

Command Default

No instance ID specified displays IS-IS adjacencies for all the IS-IS instances.

Both Level 1 and Level 2 are configured if no level is specified.
Displays all types of route calculation (not just fspf, ispf and prc).

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |
| | Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | isis | read |

Examples The following is sample output from the **show isis spf-log** command:

```
RP/0/RSP0/CPU0:router# show isis spf-log

IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
      Time Total Trig
Timestamp  Type (ms)  Nodes Count First Trigger LSP Triggers
-----
--- Thurs Aug 19 2004 ---
12:00:50.787 FSPF 1 1 3 ensoft-grs7.00-00 LSPHEADER TLVCODE
12:00:52.846 FSPF 1 1 1 ensoft-grs7.00-00 LSPHEADER
12:00:56.049 FSPF 1 1 1 ensoft-grs7.00-00 TLVCODE
12:01:02.620 FSPF 1 1 2 ensoft-grs7.00-00 NEWADJ LINKTLV

IS-IS 1 Level 1 IPv4 Unicast Route Calculation Log
      Time Total Trig
Timestamp  Type (ms)  Nodes Count First Trigger LSP Triggers
-----
--- Mon Aug 19 2004 ---
12:00:50.790 FSPF 0 1 4 ensoft-grs7.00-00 LSPHEADER TLVCODE
12:00:54.043 FSPF 1 1 2 ensoft-grs7.00-00 NEWADJ LSPHEADER
12:00:55.922 FSPF 1 2 1 ensoft-grs7.00-00 NEWLSPO
12:00:56.724 FSPF 1 13 1 ensoft-grs7.00-00 NEWLSPO
```

This table describes the significant fields shown in the display.

Table 80: show isis spf-log ipv4 unicast Field Descriptions

| Field | Description |
|-------------------|---|
| Level | IS-IS level of the router. |
| Timestamp | Time when the SPF calculation started. |
| Duration | Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time. |
| Nodes | Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run. |
| Trig Count | Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the spf-interval command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see <i>List of Triggers</i> . |
| First Trigger LSP | LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered. |
| Triggers | List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see <i>List of Triggers</i> . |

This table lists triggers of a full SPF calculation.

Table 81: List of Triggers

| Trigger | Description |
|---------------|--|
| PERIODIC | Runs a full SPF calculation every 15 minutes. |
| NEWLEVEL | Configured new level (using is-type) on this router. |
| RTCLEARED | Cleared IS-IS topology on the router. |
| MAXPATHCHANGE | Changed IP maximum parallel path. |
| NEWMETRIC | Changed link metric. |

| Trigger | Description |
|------------|--|
| ATTACHFLAG | Changed Level 2 Attach bit. |
| ADMINDIST | Configured another administrative distance for the IS-IS instance on this router. |
| NEWADJ | Created a new adjacency to another router. |
| DELADJ | Deleted adjacency. |
| BACKUP | Installed backup route. |
| SEEDISPF | Seed incremental SPF. |
| NEXTHOP | Changed IP next-hop address. |
| NEWLSP0 | New LSP 0 appeared in the topology. |
| LSPEXPIRED | Some LSP in the link-state database (LSDB) has expired. |
| LSPHEADER | Changed important LSP header fields. |
| TLVCODE | Type, length, and value (TLV) objects code mismatch, indicating that different TLV objects are included in the newest version of an LSP. |
| LINKTV | Changed Link TLV content. |
| PREFIXTLV | Changed Prefix TLV content. |
| AREAADDRTL | Changed Area address TLV content. |
| IP ADDRTL | Changed IP address TLV content. |
| TUNNEL | Changed RRR tunnel. |

The following is sample output from the **show isis spf-log** command with the **first** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis spf-log first 2

IISIS isp Level 1 IPv4 Unicast Route Calculation Log
      Time Total Trig
Timestamp Type (ms) Nodes Count First Trigger LSP Triggers
Mon Aug 16 2004
19:25:35.140 FSPF 1 1 1 12a5.00-00 NEWLSP0
19:25:35.646 FSPF 1 1 1 NEWADJ

IISIS isp Level 2 IPv4 Unicast Route Calculation Log
      Time Total Trig
Timestamp Type (ms) Nodes Count First Trigger LSP Triggers
```

```

Mon Aug 16 2004
19:25:35.139 FSPF 1 1 1 12a5.00-00 NEWLSP0
19:25:35.347 FSPF 1 1 2 12a5.00-00 NEWSADJ TLVCODE

```

This table describes the significant fields shown in the display.

Table 82: show isis spf-log first Field Descriptions

| Field | Description |
|-------------------|---|
| Level | IS-IS level of the router. |
| Timestamp | Time at which the SPF calculation started. |
| Type | Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC). |
| Time (ms) | Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time. |
| Nodes | Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run. |
| Trig Count | Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the spf-interval command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see <i>List of Triggers</i> . |
| First Trigger LSP | LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered. |
| Triggers | List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see <i>List of Triggers</i> . |

The following is sample output from the **show isis spf-log** command with the **detail** keyword specified:

```

RP/0/RSP0/CPU0:router# show isis spf-log detail

ISIS isp Level 1 IPv4 Unicast Route Calculation Log
Time Total Trig
Timestamp Type (ms) Nodes Count First Trigger LSP Triggers

```

```

Mon Aug 16 2004
19:25:35.140 FSPF 1 1 1 12a5.00-00 NEWLSP0
Delay: 51ms (since first trigger)
SPT Calculation
  CPU Time: 0ms
  Real Time: 0ms
Prefix Updates
  CPU Time: 1ms
  Real Time: 1ms
New LSP Arrivals: 0
Next Wait Interval: 200ms

Results
Reach Unreach Total
Nodes: 1 0 1
Prefixes (Items)
  Critical Priority: 0 0 0
  High Priority: 0 0 0
  Medium Priority 0 0 0
  Low Priority 0 0 0
  All Priorities 0 0 0
Prefixes (Routes)
  Critical Priority: 0 - 0
  High Priority: 0 - 0
  Medium Priority 0 - 0
  Low Priority: 0 - 0
  All Priorities 0 - 0

```

This table describes the significant fields shown in the display.

Table 83: show isis spf-log detail Field Descriptions

| Field | Description |
|------------|---|
| Level | IS-IS level of the router. |
| Timestamp | Time at which the SPF calculation started. |
| Type | Type of route calculation. The possible types are incremental SPF (iSPF), full SPF (FSPF), or partial route calculation (PRC). |
| Time (ms) | Number of milliseconds taken to complete this SPF run. Elapsed time is wall clock time, not CPU time. |
| Nodes | Number of routers and pseudonodes (LANs) that make up the topology calculated in this SPF run. |
| Trig Count | Number of events that triggered this SPF run. When there is a topology change, often multiple link-state packets (LSPs) are received in a short time. Depending on the configuration of the spf-interval command, a router may wait for a fixed period of time before running a router calculation. This count denotes the number of triggering events that occurred while the router was waiting to run the calculation. For a full description of the triggering events, see <i>List of Triggers</i> . |

| Field | Description |
|--------------------|---|
| First Trigger LSP | LSP ID stored by the router whenever a full SPF calculation is triggered by the arrival of a new LSP. The LSP ID can suggest the source of routing instability in an area. If multiple LSPs are causing an SPF run, only the LSP ID of the first received LSP is remembered. |
| Triggers | List of all reasons that triggered a full SPF calculation. For a list of possible triggers, see <i>List of Triggers</i> . |
| Delay | Two different delays exist: <ol style="list-style-type: none"> 1 The delay between the time when the route calculation was first triggered and the time when it was run. 2 The delay between the end of the last route calculation and the start of this one. This is used to verify that the SPF-interval timers are working correctly, and is only reported for calculations after the first delay. |
| CPU Time | Two different CPU times exist: <ol style="list-style-type: none"> 1 CPU time (in milliseconds) taken to calculate the shortest path tree (SPT). 2 CPU time (in milliseconds) taken to perform the prefix updates. |
| Real Time | Two different real times exist: <ol style="list-style-type: none"> 1 Real time (in milliseconds) taken to calculate the shortest path tree (SPT). 2 Real time (in milliseconds) taken to perform the prefix updates. |
| New LSP Arrivals | Number of LSP arrivals since the start of this route calculation. |
| Next Wait Interval | Enforced delay until the next route calculation can be run, based on the spf-interval command configuration. |
| Reach | Number of reachable nodes or prefixes. |
| Unreach | Number of unreachable nodes or prefixes. |
| Total | Total number of nodes or prefixes at various priorities. |

Related Commands

| Command | Description |
|--|--|
| spf-interval , on page 913 | Sets IS-IS throttling of shortest path first (SPF) calculations. |

show isis statistics

To display Intermediate System-to-Intermediate System (IS-IS) traffic counters, use the **show isis statistics** command in EXEC mode.

show isis [**instance** *instance-id*] **statistics** [*type interface-path-id*]

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS traffic statistics for the specified IS-IS instance only. <ul style="list-style-type: none"> The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. <p>Note Use the show interfaces command to see a list of all interfaces currently configured on the router.</p> <p>For more information about the syntax for the router, use the question mark (?) online help function.</p> |

Command Default

No instance ID specified displays IS-IS traffic statistics for all the IS-IS instances. IS-IS traffic statistics are displayed for all interfaces.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **show isis statistics** command displays IS-IS traffic counters for the specified interface or all traffic counters if no interface is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis statistics** command that shows all traffic counters:

```

RP/0/RSP0/CPU0:router#show isis statistics
IS-IS isp statistics:
  Fast PSNP cache (hits/tries): 164115/301454
  Fast CSNP cache (hits/tries): 41828/43302
  Fast CSNP cache updates: 2750
  LSP checksum errors received: 0
  LSP Dropped: 1441
  SNP Dropped: 1958
  UPD Max Queue size: 2431
Average transmit times and rate:
  Hello:      0 s,      987947 ns,      4/s
  CSNP:      0 s,     1452987 ns,      0/s
  PSNP:      0 s,     1331690 ns,      0/s
  LSP:       0 s,     1530018 ns,      1/s
Average process times and rate:
  Hello:      0 s,      874584 ns,      41/s
  CSNP:      0 s,      917925 ns,      29/s
  PSNP:      0 s,     1405458 ns,      0/s
  LSP:       0 s,     4352850 ns,      0/s
Level-1:
  LSPs sourced (new/refresh): 3376/2754
  Level-1:LSPs sourced (new/refresh): 3376/2754IPv4 UnicastSPF calculations
: 520ISPF calculations      : 0
  Next-hop Calculations      : 0
  Partial Route Calculations : 0
  IPv6 Unicast
  SPF calculations           : 527
  ISPF calculations          : 0
  Next Hop Calculations      : 13
  Partial Route Calculations : 1
  Level-2:
  LSPs sourced (new/refresh): 4255/3332
  IPv4 Unicast
  SPF calculations           : 432
  ISPF calculations          : 0
  Next Hop Calculations      : 8
  LSPs sourced (new/refresh): 4255/3332LSPs sourced (new/refresh): 4255/3332
  IPFRR Parallel calculations: 0
  IPv4 IPv6 Unicast
  SPF calculations           : 432
  ISPF calculations          : 0
  Next-hop Calculations      : 8
  Partial Route Calculations : 0
Interface GigabitEthernet0/1/0/1.1:
Level-1 Hellos (sent/rcvd): 22398/25633
Level-1 DR Elections      : 66
Level-1 LSPs (sent/rcvd) : 246/7077
Level-1 CSNPs (sent/rcvd) : 0/33269
Level-1 PSNPs (sent/rcvd) : 22/0
Level-1 LSP Flooding Duplicates : 25129
Level-2 Hellos (sent/rcvd): 22393/67043
Level-2 DR Elections      : 55
Level-2 LSPs (sent/rcvd) : 265/437
Level-2 CSNPs (sent/rcvd) : 0/86750
Level-2 PSNPs (sent/rcvd) : 0/0
Level-2 LSP Flooding Duplicates : 78690

```

This table describes the significant fields shown in the display.

Table 84: show isis statistics Field Descriptions

| Field | Description |
|---------------------------------|--|
| Fast PSNP cache (hits/tries) | Number of successful lookups (hits) along with the number of lookup attempts (tries). To save time or processing power when receiving multiple copies of the same LSP, IS-IS attempts to look up incoming LSPs to see if they have been received recently. |
| Fast CSNP cache (hits/tries) | Number of successful lookups (hits) along with the number of lookup attempts (tries). To reduce CSNP construction time, IS-IS maintains a cache of CSNPs and attempts to look up CSNP in this cache before transmission on the interface. |
| Fast CSNP cache updates | Number of times the CSNP cache has been updated since the last clearing of statistics. The cache is updated on LSP addition or removal from the database. |
| LSP checksum errors received | Number of internal checksum errors received in LSPs. |
| IIH (LSP/SNP) dropped | Number of hello, LSP, and SNP messages dropped. |
| IIH (UPD) Max Queue size | Maximum number of queued packets. |
| Average transmit times and rate | Average time taken to transmit the pdu type across all interfaces and the corresponding rate at which the pdu type is being transmitted. |
| Average process times and rate | Average time taken to process an incoming pdu type across all interfaces and the corresponding rate at which the pdu type is being received. |
| LSPs sourced (new/refresh) | Number of LSPs this IS-IS instance has created or refreshed. To find more details on these LSPs, use the show isis lsp-log command. |
| SPF calculations | Number of shortest path first (SPF) calculations. SPF calculations are performed only when the topology changes. They are not performed when external routes change. The interval at which SPF calculations are performed is configured using the spf-interval command. |
| iSPF calculations | Number of incremental shortest path first (iSPF) calculations. iSPF calculations are performed only when ISPF has been configured in the isis address family configuration submode. |

| Field | Description |
|---|--|
| Partial Route Calculations | Number of partial route calculations (PRCs). PRCs are processor intensive. Therefore, it may be useful to limit their number, especially how often a PRC is done, especially on slower networking devices. Increasing the PRC interval reduces the processor load on the router, but might slow the rate of convergence. The interval at which PRC calculations are performed is configured using the spf-interval command. |
| Level-(1/2) (LSPs/CSNPs/PSNPs/Hellos) (sent/rcvd) | Number of LSPs, Complete Sequence Number Packets (CSNPs), Partial Sequence Number Packets (PSNPs), and hello packets sent or received on this interface. |
| PTP Hellos (sent/rcvd) | Point-to-point (PTP) hellos sent and received. |
| LSP Retransmissions | Total number of retransmissions on each IS-IS LSP on a point-to-point interface. The LSP retransmission interval can be configured using the retransmit-throttle-interval command. |
| Level-(1.2) DRElections | Total number of Designated Intermediate System elections that have taken place. These counts are maintained on an individual level basis. |
| LSP Flooding Duplicates | Number of duplicate LSPs filtered from flooding to the neighbor. In case of parallel interfaces to the same neighbor, IS-IS optimizes the flooding by avoiding sending the same LSP copy on other interfaces. |

Related Commands

| Command | Description |
|---|---|
| show isis spf-log , on page 890 | Displays how often and why the router has run a full SPF calculation. |
| spf-interval , on page 913 | Sets IS-IS throttling of shortest path first (SPF) calculations. |

show isis topology

To display a list of connected Intermediate System-to-Intermediate System (IS-IS) routers in all areas, use the **show isis topology** command in EXEC mode.

```
show isis [instance instance-id] [[ipv4|ipv6] afi-all] [unicast|multicast [topology {all|topo-name}]]
safi-all]] summary| level {1|2} [multicast-intact] [systemid system-id] [detail]
```

Syntax Description

| | |
|------------------------------------|---|
| instance <i>instance-id</i> | (Optional) Displays the IS-IS topology for the specified IS-IS instance only. • The <i>instance-id</i> argument is the instance identifier (alphanumeric) defined by the router isis command. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| afi-all | (Optional) Specifies all address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| safi-all | (Optional) Specifies all secondary address prefixes. |
| summary | (Optional) Displays a brief list of the IS-IS topology. |
| level { 1 2 } | (Optional) Displays the IS-IS link-state topology for Level 1 or Level 2 independently. |
| multicast-intact | (Optional) Displays multicast intact information on the IS-IS topology. |
| systemid <i>system-id</i> | (Optional) Displays the information for the specified router only. |
| detail | (Optional) Displays detailed information on the IS-IS topology. |

Command Default

No instance ID specified displays a list of connected routers in all areas for all the IS-IS instances. Both Level 1 and Level 2 is configured if no level is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support for IPv6 was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show isis topology** command to verify the presence and connectivity among all routers in all areas.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |

Examples

The following is sample output from the **show isis topology** command:

```
RP/0/RSP0/CPU0:router# show isis topology

IS-IS isp paths to (Level-1) routers
 System Id      Metric  Next-hop Interface  SNPA
 ensoft-5       10     ensoft-5   PO0/4/0/1    *PtoP*
 ensoft-5       10     ensoft-5   Gi0/5/0/0    0003.6cff.0680
 ensoft-11      --

IS-IS isp paths to (Level-2) routers
 System Id      Metric  Next-hop Interface  SNPA
 ensoft-5       10     ensoft-5   PO0/4/0/1    *PtoP*
 ensoft-5       10     ensoft-5   Gi0/5/0/0    0003.6cff.0680
 ensoft-11      --
```

This table describes the significant fields shown in the display.

Table 85: show isis topology ipv4 unicast Field Descriptions

| Field | Description |
|-----------|--|
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |

| Field | Description |
|-----------|--|
| Metric | Metric assigned to the link and used to calculate the cost from each router using the links in the network to other destinations. Range is 1 to 16777214. Default is 1 to 63 for narrow metric and 1 to 16777214 for wide metric. 0 is set internally if no metric has been specified by the user. |
| Next-hop | Address of the next-hop. |
| Interface | Interface used to reach the neighbor. |
| SNPA | Data-link address (also known as the Subnetwork Point of Attachment [SNPA]) of the neighbor. |

The following is sample output from the **show isis topology** command with the **summary** keyword specified:

```
RP/0/RSP0/CPU0:router# show isis topology summary
IS-IS 10 IS Topology Summary IPv4 Unicast
          L1
          Reach  UnReach  Total
-----
Router nodes:    1      1      2
Pseudo nodes:   0      0      0
          L2
          Reach  UnReach  Total
-----
Total nodes:    1      1      2
```

This table describes the significant fields shown in the display.

Table 86: show isis topology summary Field Descriptions

| Field | Description |
|---------|---|
| L1/L2 | IS-IS level of the router. |
| Reach | Number of router nodes or pseudonodes that are reachable. |
| UnReach | Number of router nodes or pseudonodes that are unreachable. |
| Total | Total number of reachable and unreachable nodes. |

show protocols (IS-IS)

To group a number of protocol show commands according to the specified address family, use the **show protocols** command in EXEC mode.

```
show protocols [afi-all| ipv4| ipv6] [all| protocol]
```

Syntax Description

| | |
|-----------------|---|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies an IPv4 address family. |
| ipv6 | (Optional) Specifies an IPv6 address family. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • isis • ospf • rip • eigrp <p>For the IPv6 address family, the options are:</p> <ul style="list-style-type: none"> • bgp • isis • ospfv3 |

Command Default

If no address family is specified, the default is IPv4.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support for IPv6 was added |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If IPv6 is enabled on an IS-IS instance, the instance is displayed in the **show protocols ipv6** command output. IPv4 IS-IS instances are displayed in the **show protocols ipv4** command output.

When using the **show protocols** command with the **ipv6** or **ipv4** keyword, you get all routing instances in that particular address family—not only IS-IS instances.

Task ID

| Task ID | Operations |
|---------|------------|
| isis | read |
| rib | read |

Examples

The following example shows the output for the **show protocols** command :

```
RP/0/RSP0/CPU0:router# show protocols ipv4

IS-IS Router: uut
System Id: 0000.0000.12a8
IS Levels: level-1-2
Manual area address(es):
 49.1515.1515
Routing for area address(es):
 49.1515.1515
Non-stop forwarding: Disabled
Most recent startup mode: Cold Restart
Topologies supported by IS-IS:
 IPv4 Unicast
   Level-1
     Metric style (generate/accept): Narrow/Narrow
     ISPF status: Disabled
   Level-2
     Metric style (generate/accept): Narrow/Narrow
     ISPF status: Disabled
   Redistributing:
     static
     Distance: 115
 IPv6 Unicast
   Level-1
     ISPF status: Disabled
   Level-2
     ISPF status: Disabled
   No protocols redistributed
   Distance: 45
Interfaces supported by IS-IS:
 GigabitEthernet 0/6/0/0 is running actively (active in configuration)
```

This table describes the significant fields shown in the display.

Table 87: show protocols ipv4 Field Descriptions

| Field | Description |
|-------------------------------|--|
| System ID | Dynamic hostname of the system. The hostname is specified using the hostname command. If the dynamic hostname is not known or hostname dynamic disable command has been executed, the 6-octet system ID is used. |
| IS Levels | IS-IS level of the router. |
| Manual area address(es) | Area addresses configured manually on the originating router. |
| Routing for area address(es) | Area addresses for which this router provides the routing. |
| Non-stop forwarding | Status and name of NSF. |
| Most recent startup mode | Mode in which the most recent startup was performed. |
| Topologies supported by IS-IS | Address and subaddress family IS-IS are configured. |
| Metric style | Type, length, and value (TLV) objects accepted by IS-IS. To configure this value, see the metric-style narrow, on page 784 , metric-style transition, on page 786 , or metric-style wide, on page 788 command. |
| ISPF status | State of iSPF configuration for this IS-IS instance. Four states exist: <ul style="list-style-type: none"> • Disabled if iSPF has not been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Dormant if iSPF has been configured but is awaiting initial convergence before initializing. • Awake if iSPF has been configured but is awaiting a full SPF to compile the topology for use by the iSPF algorithm. • Active if IS-IS is ready to consider using the iSPF algorithm whenever a new route calculation needs to be run. |
| Redistributing | IS-IS is configured to redistribute IP static routes into Level 1 or Level 2. The redistribute command is used to configure redistribution. |
| Distance | Administrative distance. |

| Field | Description |
|-------------------------------|--|
| Interfaces supported by IS-IS | Interfaces and their states currently supported by IS-IS. Both operational and configuration status are displayed. |

The following example shows how to disable the IPv4 address family, with no output shown for IS-IS IPv4 instances from the **show protocols ipv4** command:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router isis uut
RP/0/RSP0/CPU0:router(config-isis)# no address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis)# commit

RP/0/RSP0/CPU0:router# show protocols ipv4
```

Related Commands

| Command | Description |
|--|---|
| metric-style narrow, on page 784 | Configures the IS-IS software to generate and accept old-style type, length, and value (TLV) objects. |
| metric-style transition, on page 786 | Configures the IS-IS software to generate and accept both old-style and new-style type length, and value (TLV) objects. |
| metric-style wide, on page 788 | Configures the IS-IS software to generate and accept only new-style type, length, and value (TLV) objects. |
| redistribute (IS-IS), on page 819 | Redistributes routes from one IS-IS instance into another instance. |

shutdown (IS-IS)

To disable the Intermediate System-to-Intermediate System (IS-IS) protocol on a particular interface, use the **shutdown** command in interface configuration mode. To re-enable the IS-IS protocol, use the **no** form of this command.

shutdown

no shutdown

Command Default IS-IS protocol is enabled.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example disables the IS-IS protocol on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# shutdown
```

single-topology

To configure the link topology for IP Version 4 (IPv4) when IP Version 6 (IPv6) is configured, use the **single-topology** command in address family configuration mode. To remove the **single-topology** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

single-topology

no single-topology

Command Default

Performs in multitopology mode in which independent topologies for IPv4 and IPv6 are running in a single area or domain.

Command Modes

IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **single-topology** command to allow Intermediate System-to-Intermediate System (IS-IS) for IPv6 to be configured on interfaces along with an IPv4 network protocol. All interfaces must be configured with the identical set of network protocols, and all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer protocols on all interfaces.

When single-topology support for IPv6 is being used, only old-style type, length, and value (TLV) objects may be used and a single shortest path (SPF) individual level is used to compute IPv4 (if configured) and IPv6 routes. The use of a single SPF means that both IPv4 IS-IS and IPv6 IS-IS routing protocols must share a network topology.

To allow link information to be shared between IPv4 and IPv6, you must configure the **single-topology** command for an address family. In single-topology IPv6 mode, the configured metric is always the same for both IPv4 and IPv6.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to enable single-topology mode for IPv6:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# net 49.0000.0000.0001.00  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv6 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# single-topology
```

snmp-server traps isis

snmp-server traps isis {all| *traps set*}

no snmp-server traps isis {all| *traps set*}

Command Default

Command Modes

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Examples

```
RP/0/RSP0/CPU0:router(config)# snmp-server traps isis

adjacency-change          isisAdjacencyChange
all                        Enable all IS-IS traps
area-mismatch             isisAreaMismatch
attempt-to-exceed-max-sequence isisAttemptToExceedMaxSequence
authentication-failure    isisAuthenticationFailure
authentication-type-failure isisAuthenticationTypeFailure
corrupted-lsp-detected    isisCorruptedLSPDetected
database-overload         isisDatabaseOverload
id-len-mismatch           isisIDLenMismatch
lsp-error-detected        isisLSPErrorDetected
lsp-too-large-to-propagate isisLSPTooLargeToPropagate
manual-address-drops      isisManualAddressDrops
max-area-addresses-mismatch isisMaxAreaAddressesMismatch
orig-lsp-buff-size-mismatch isisOrigLSPBuffSizeMismatch
own-lsp-purge             isisOwnLSPPurge
protocols-supported-mismatch isisProtocolsSupportedMismatch
rejected-adjacency        isisRejectedAdjacency
sequence-number-skip      isisSequenceNumberSkip
version-skew              isisVersionSkew

RP/0/RSP0/CPU0:router(config)#snmp-server traps isis all

RP/0/RSP0/CPU0:router(config)# snmp-server traps isis area-mismatch
lsp-error-detected
```

spf-interval

To customize IS-IS throttling of shortest path first (SPF) calculations, use the **spf-interval** command in address family configuration mode. To restore default values, use the **no** form of this command.

```
spf-interval [initial-wait initial| secondary-wait secondary| maximum-wait maximum] ... [level {1|2}]
no spf-interval [[initial-wait initial| secondary-wait secondary| maximum-wait maximum] ...] [level {1|2}]
```

Syntax Description

| | |
|--|--|
| initial-wait <i>initial</i> | Initial SPF calculation delay (in milliseconds) after a topology change. Range is 0 to 120000. |
| secondary-wait <i>secondary</i> | Hold time between the first and second SPF calculations (in milliseconds). Range is 0 to 120000. |
| maximum-wait <i>maximum</i> | Maximum interval (in milliseconds) between two consecutive SPF calculations. Range is 0 to 120000. |
| level { 1 2 } | (Optional) Enables the SPF interval configuration for Level 1 or Level 2 independently. |

Command Default

```
initial-wait initial : 50 milliseconds
secondary-wait secondary : 200 milliseconds
maximum-wait maximum : 5000 milliseconds
```

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

SPF calculations are performed only when the topology changes. They are not performed when external routes change.

Use the **spf-interval** command to control how often the software can perform the SPF calculation. The SPF calculation is processor intensive. Therefore, it may be useful to limit how often this calculation is done,

especially when the area is large and the topology changes often. Increasing the SPF interval reduces the processor load of the router, but potentially slows the rate of convergence.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| isis | read, write |

Examples

The following example shows how to set the initial SPF calculation delay to 10 milliseconds and the maximum interval between two consecutive SPF calculations to 5000 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router isis isp  
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)# spf-interval initial-wait 10 maximum-wait 5000
```

spf prefix-priority (IS-IS)

To assign a priority to an ISIS prefix for customizing the RIB update sequence, use the **spf prefix-priority** command in address family configuration mode. To restore default values, use the **no** form of this command.

```
spf prefix-priority [level {1|2}] {critical|high|medium} {access-list-name|tag tag}
```

```
no spf prefix-priority [level {1|2}] {critical|high|medium} [access-list-name|tag tag]
```

Syntax Description

| | |
|-------------------------|---|
| level { 1 2 } | (Optional) Enables the assignment of a priority to Level 1 or Level 2 independently. |
| critical | Assigns a critical priority. |
| high | Assigns a high priority. |
| medium | Assigns a medium priority. |
| <i>access-list-name</i> | Name of an access list. |
| tag <i>tag</i> | Specifies a tag to indicate priority. The <i>tag</i> argument range is 1 to 4294967295. |

Command Default

By default, IPv4 prefixes with a length of 32 and IPv6 prefixes with a length of 128 are given medium priority. The remaining prefixes are given low priority.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **spf prefix-priority** command to change the sequence of prefix updates to the RIB after an SPF is run. ISIS installs prefixes in the RIB according to the following priority order:

Critical > High > Medium > Low

The **spf prefix-priority** command supports prefix lists for the first three priorities. The unmatched prefixes are updated with low priority.

If a **spf prefix-priority** is specified, the default behavior of prioritizing either length 32 or 128 prefixes for IPv4 or IPv6, respectively, as **medium** is disabled.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to set the prefix priorities:

```
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-critical-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 32
!
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-med-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 29
!
RP/0/RSP0/CPU0:router(config)# ipv4 prefix-list isis-high-acl
RP/0/RSP0/CPU0:router(config-ipv4_pfx)# 10 permit 0.0.0.0/0 eq 30
!
RP/0/RSP0/CPU0:router(config)# router isis ring
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority critical isis-critical-acl
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority high isis-high-acl
RP/0/RSP0/CPU0:router(config-isis-af)# spf prefix-priority medium isis-med-acl
```

summary-prefix (IS-IS)

To create aggregate addresses for the Intermediate System-to-Intermediate System (IS-IS) protocol, use the **summary-prefix** command in address family configuration mode. To restore the default behavior, use the **no** form of this command.

Syntax Description

| | |
|-------------------------------|--|
| address | Summary address designated for a range of IPv4 addresses. The <i>address</i> argument must be in four-part, dotted-decimal notation. |
| / <i>prefix-length</i> | Length of the IPv4 or IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. |
| ipv6-prefix | Summary prefix designated for a range of IPv6 prefixes. The <i>ipv6-prefix</i> argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons. |
| level { 1 2 } | (Optional) Redistributes routes into Level 1 or Level 2 and summarizes them with the configured address and mask value. |
| tag <i>tag</i> | Sets a tag value. The value range is 1- 4294967295. |

Command Default

All redistributed routes are advertised individually.
Both Level 1 and Level 2 are configured if no level is specified.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Tag keyword and IPv6 support was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Multiple groups of addresses can be summarized for a given level. Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the more-specific routes. Use the **summary-prefix** command to help reduce the size of the routing table.

This command also reduces the size of the link-state packets (LSPs) and thus the link-state database. It also helps ensure stability, because a summary advertisement depends on many more specific routes. If one more-specific route flaps, in most cases, this flap does not cause a flap of the summary advertisement.

The drawback of summary addresses is that other routes might have less information to calculate the most optimal routing table for all individual destinations.

**Note**

When IS-IS advertises a summary prefix, it automatically inserts the summary prefix into the IP routing table but labels it as a “discard” route entry. Any packet that matches the entry is discarded to prevent routing loops. When IS-IS stops advertising the summary prefix, the routing table entry is removed.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to redistribute Open Shortest Path First (OSPF) routes into IS-IS:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 ipv6 unicast
RP/0/RSP0/CPU0:router(config-isis-af)# redistribute ospf 2 level-2
RP/0/RSP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10 level-2
RP/0/RSP0/CPU0:router(config-isis-af)# summary-prefix 10.10.10.10
```

suppressed

To allow an IS-IS interface to participate in forming adjacencies without advertising connected prefixes in the system link-state packets (LSPs), use the **suppressed** command in interface configuration mode. To enable advertising connected prefixes, use the **no** form of this command.

suppressed

no suppressed

Command Default Interface is active.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **suppressed** command to reduce the number of routes that IS-IS has to maintain, improving convergence times after an isolated failure. Improvement is noticeable if the command is used widely throughout the network. Other routers in the domain do not install routes to the affected connected prefixes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | isis | read, write |

Examples The following example shows how to disable the advertisement of connected prefixes on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet /1/0/1
RP/0/RSP0/CPU0:router(config-isis-if)# suppressed
```

| Related Commands | Command | Description |
|------------------|--|--|
| | passive (IS-IS), on page 813 | Suppresses S-IS packets on an interface. |

■ suppressed

tag (IS-IS)

To associate and advertise a tag with the prefix of an IS-IS interface, use the **tag** command in interface address family configuration mode. To restore the default behavior, use the **no** form of this command.

tag *tag*

no tag [*tag*]

Syntax Description

tag Interface tag. Range is 1 to 4294967295.

Command Default

Default is that no tag is associated and advertised.

Command Modes

Interface address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to associate and advertise an interface tag:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-isis-if)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-isis-if-af)# tag 234
```

Related Commands

| Command | Description |
|---|---|
| spf prefix-priority (IS-IS) , on page 915 | Assigns a priority to an ISIS prefix for customizing the RIB update sequence. |

topology-id

To differentiate one topology in the domain from another while configuring a multicast routing table, use the **topology-id** command in Intermediate System-to-Intermediate System (IS-IS) address family configuration submode. To disable the topology use the **no** form of the command.

topology-id *isis-multicast-topology-id-number*

no topology-id *isis-multicast-topology-id-number*

Syntax Description

| | |
|--|--|
| <i>isis-multicast-topology-id-number</i> | ID number for a specific IS-IS multicast topology. Range is 6 to 4095. |
|--|--|

Command Default

No topology is associated with a routing table by default.

Command Modes

IS-IS address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to differentiate a topology from another in the multicast routing table in IS-IS routing:

```
RP/0/RSP0/CPU0:router(config)# router isis isp
RP/0/RSP0/CPU0:router(config-isis)# address-family ipv4 multicast topology green
RP/0/RSP0/CPU0:router(config-isis-af)# topology-id 2666
```

Related Commands

| Command | Description |
|---|---|
| address-family multicast topology (IS-IS) , on page 709 | Used in conjunction with the topology-id command, enables a multicast topology globally when configuring Intermediate System-to-Intermediate System (IS-IS) routing. |

trace (IS-IS)

To set the the IS-IS buffer size, use the **trace** command in router configuration mode. To return to the default value, use the **no** form of this command.

trace [**detailed**| **severe**| **standard**] *max-trace-entries*

no trace [**detailed**| **severe**| **standard**]

Syntax Description

| | |
|--------------------------|--|
| detailed | Specifies the buffer size for detailed traces. Range is |
| severe | Specifies the buffer size for severe traces. Range is |
| standard | Specifies the buffer size for standard traces. Range is |
| <i>max-trace-entries</i> | Sets the maximum number of trace entries. Range is 1-20000 |

Command Default

None

Command Modes

Router IS-IS configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| isis | read, write |

Examples

The following example shows how to set the isis buffer size for severe traces to 1200:

```
RP/0/RSP0/CPU0:router(config)#router isis isp  
RP/0/RSP0/CPU0:router(config-isis)#trace sever 1200
```



OSPF Commands

This module describes the commands used to configure and monitor the Open Shortest Path First (OSPF) routing protocol.

For detailed information about OSPF concepts, configuration tasks, and examples, see the *Implementing OSPF on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family \(OSPF\)](#), page 931
- [adjacency stagger](#), page 932
- [area \(OSPF\)](#), page 934
- [authentication \(OSPF\)](#), page 936
- [authentication-key \(OSPF\)](#), page 939
- [auto-cost \(OSPF\)](#), page 941
- [capability opaque disable](#), page 943
- [capability vrf-lite \(OSPFv2\)](#), page 944
- [clear ospf process](#), page 946
- [clear ospf redistribution](#), page 948
- [clear ospf routes](#), page 950
- [clear ospf statistics](#), page 952
- [cost \(OSPF\)](#), page 954
- [cost-fallback \(OSPF\)](#), page 956
- [database-filter all out \(OSPF\)](#), page 958
- [dead-interval \(OSPF\)](#), page 960
- [default-cost \(OSPF\)](#), page 962
- [default-information originate \(OSPF\)](#), page 964
- [default-metric \(OSPF\)](#), page 966
- [demand-circuit \(OSPF\)](#), page 968

- [disable-dn-bit-check](#), page 970
- [distance \(OSPF\)](#), page 971
- [distance ospf](#), page 974
- [distribute-list](#), page 976
- [domain-id \(OSPF\)](#), page 978
- [domain-tag](#), page 980
- [fast-reroute \(OSPFv2\)](#), page 981
- [fast-reroute per-link exclude interface](#), page 983
- [fast-reroute per-prefix exclude interface \(OSPFv2\)](#), page 985
- [fast-reroute per-prefix lfa-candidate \(OSPFv2\)](#), page 987
- [fast-reroute per-prefix remote-lfa \(OSPFv2\)](#), page 988
- [fast-reroute per-prefix use-candidate-only \(OSPFv2\)](#), page 990
- [flood-reduction \(OSPF\)](#), page 991
- [hello-interval \(OSPF\)](#), page 993
- [ignore lsa mospf](#), page 995
- [interface \(OSPF\)](#), page 996
- [log adjacency changes \(OSPF\)](#), page 998
- [loopback stub-network](#), page 999
- [max-lsa](#), page 1001
- [max-metric](#), page 1004
- [maximum interfaces \(OSPF\)](#), page 1007
- [maximum paths \(OSPF\)](#), page 1009
- [maximum redistributed-prefixes \(OSPF\)](#), page 1011
- [message-digest-key](#), page 1013
- [mpls ldp auto-config \(OSPF\)](#), page 1016
- [mpls ldp sync \(OSPF\)](#), page 1017
- [mpls traffic-eng \(OSPF\)](#), page 1018
- [mpls traffic-eng igp-intact \(OSPF\)](#), page 1020
- [mpls traffic-eng multicast-intact \(OSPF\)](#), page 1022
- [mpls traffic-eng router-id \(OSPF\)](#), page 1024
- [mtu-ignore \(OSPF\)](#), page 1026
- [multi-area-interface](#), page 1028
- [neighbor \(OSPF\)](#), page 1030

- [neighbor database-filter all out](#), page 1033
- [network \(OSPF\)](#), page 1035
- [nsf \(OSPF\)](#), page 1037
- [nsf flush-delay-time \(OSPF\)](#), page 1039
- [nsf interval \(OSPF\)](#), page 1040
- [nsf lifetime \(OSPF\)](#), page 1042
- [nsr \(OSPF\)](#), page 1044
- [nssa \(OSPF\)](#), page 1045
- [ospf name-lookup](#), page 1047
- [packet-size \(OSPF\)](#), page 1048
- [passive \(OSPF\)](#), page 1050
- [priority \(OSPF\)](#), page 1052
- [protocol shutdown](#), page 1054
- [queue dispatch flush-lsa](#), page 1055
- [queue dispatch incoming](#), page 1057
- [queue dispatch rate-limited-lsa](#), page 1059
- [queue dispatch spf-lsa-limit](#), page 1061
- [queue limit](#), page 1063
- [range \(OSPF\)](#), page 1065
- [redistribute \(OSPF\)](#), page 1067
- [retransmit-interval \(OSPF\)](#), page 1072
- [route-policy \(OSPF\)](#), page 1074
- [router-id \(OSPF\)](#), page 1076
- [router ospf](#), page 1078
- [security ttl \(OSPF\)](#), page 1080
- [sham-link](#), page 1082
- [show ospf](#), page 1084
- [show ospf border-routers](#), page 1088
- [show ospf database](#), page 1090
- [show ospf flood-list](#), page 1104
- [show ospf interface](#), page 1106
- [show ospf mpls traffic-eng](#), page 1109
- [show ospf message-queue](#), page 1114

- [show ospf neighbor](#), page 1117
- [show ospf request-list](#), page 1125
- [show ospf retransmission-list](#), page 1128
- [show ospf routes](#), page 1131
- [show ospf sham-links](#), page 1136
- [show ospf summary-prefix](#), page 1139
- [show ospf virtual-links](#), page 1141
- [show protocols \(OSPF\)](#), page 1143
- [snmp context \(OSPF\)](#), page 1146
- [snmp trap \(OSPF\)](#), page 1149
- [snmp trap rate-limit \(OSPF\)](#), page 1150
- [spf prefix-priority \(OSPFv2\)](#), page 1151
- [stub \(OSPF\)](#), page 1153
- [summary-prefix \(OSPF\)](#), page 1155
- [timers lsa group-pacing](#), page 1157
- [timers lsa min-arrival](#), page 1159
- [timers throttle lsa all \(OSPF\)](#), page 1160
- [timers throttle spf \(OSPF\)](#), page 1163
- [transmit-delay \(OSPF\)](#), page 1165
- [virtual-link \(OSPF\)](#), page 1167
- [vrf \(OSPF\)](#), page 1169

address-family (OSPF)

To enter address family configuration mode for Open Shortest Path First (OSPF), use the **address-family** command in the appropriate mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv4 [unicast]

no address-family ipv4 [unicast]

Syntax Description

| | |
|----------------|---|
| ipv4 | Specifies IP Version 4 (IPv4) address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |

Command Default

An address family is not specified.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF version 2 automatically provides routing services for IPv4 unicast topologies, so this command is redundant.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the OSPF router process with IPv4 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# address-family ipv4 unicast
```

adjacency stagger

To configure staggering of OSPF adjacency during reload, process restart, and process clear, use the **adjacency stagger** command in router configuration mode. To turn off adjacency staggering, either use the **disable** keyword or use the **no** form of this command.

adjacency stagger {**disable**| *initial-num-nbr max-num-nbr*}

no adjacency stagger

Syntax Description

| | |
|------------------------|--|
| disable | Disables adjacency staggering. |
| <i>initial-num-nbr</i> | The initial number of simultaneous neighbors allowed to form adjacency to FULL in any area to bring up to FULL after a router reload, OSPF process restart, or OSPF process clear. Range is 1-65535. Default is 2. |
| <i>max-num-nbr</i> | The subsequent number of simultaneous neighbors allowed to form adjacency, per OSPF instance, after the initial set of OSPF neighbors have become FULL. Range is 1-65535. Default is 64. |

Command Default

OSPF adjacency staggering is enabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Staggering of the OSPF adjacency during reload, process restart (without NSF or graceful-restart), and process clear reduces the overall adjacency convergence time.

Initially, allow 2 (configurable) neighbors to form adjacency to FULL per area. After the first adjacency reaches FULL, up to 64 (configurable) neighbors can form adjacency simultaneously for the OSPF instance (all areas). However, areas without any FULL adjacency is restricted by the initial area limit.



Note

Adjacency stagger and OSPF nonstop forwarding (NSF) are mutually exclusive. Adjacency stagger will not be activated if **nsf** is configured under router ospf configuration.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows how to configure adjacency stagger for a 2 neighbors initially and for a maximum of 3 neighbors:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf)# adjacency stagger 2 3
```

area (OSPF)

To configure an Open Shortest Path First (OSPF) area, use the **area** command in the appropriate mode. To terminate an OSPF area, use the **no** form of this command.

area *area-id*

no area *area-id*

Syntax Description

| | |
|----------------|--|
| <i>area-id</i> | Identifier of an OSPF area. The <i>area-id</i> argument can be specified as either a decimal value or an IP address (dotted decimal) format. Range is 0 to 4294967295. |
|----------------|--|

Command Default

No OSPF area is defined.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **area** command to explicitly configure an area. Commands configured under the area configuration mode (such as the **interface** [OSPF] and **authentication** commands), are automatically bound to that area.

To modify or remove the area, the *area-id* argument format must be the same as the format used when creating the area. Otherwise, even if the actual 32-bit value matches, the area is not matched. For example, if you create an area with an *area-id* of 10 it would not match an *area-id* of 0.0.0.10.



Note

To remove the specified area from the router configuration, use the **no area** *area-id* command. The **no area** *area-id* command removes the area and all area options, such as **authentication**, **default-cost**, **nssa**, **range**, **stub**, **virtual-link**, and **interface**.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure area 0 and GigabitEthernet interface 0/2/0/0. GigabitEthernet interface 0/2/0/0 is bound to area 0 automatically.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/0
```

authentication (OSPF)

To enable plain text, Message Digest 5 (MD5) authentication, or null authentication for an Open Shortest Path First (OSPF) interface, use the **authentication** command in the appropriate mode. To remove such authentication, use the **no** form of this command.

authentication [**message-digest** [**keychain** *keychain*]] **null**

no authentication

Syntax Description

| | |
|---------------------------------|--|
| message-digest | (Optional) Specifies that MD5 is used. |
| keychain <i>keychain</i> | (Optional) Specifies a keychain name. |
| null | (Optional) Specifies that no authentication is used. Useful for overriding password or MD5 authentication if configured for an area. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

If no keyword is specified, plain text authentication is used.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area interface configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPF interfaces that are to communicate with each other through OSPF. If you specified plain text authentication, use the **authentication-key** command to specify the plain text password.

If you enable MD5 authentication with the **message-digest** keyword, you must configure a key with the **message-digest-key** interface command.

To manage the rollover of keys and enhance MD5 authentication for OSPF, you can configure a container of keys called a keychain with each key comprising the following attributes: generate/accept time, key identification, and authentication algorithm. The keychain management feature is always enabled.

**Note**

Changes to the system clock will impact the validity of the keys in the existing configuration.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set authentication for areas 0 and 1 of OSPF routing process 201. Authentication keys are also provided.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospf)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# authentication
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# authentication-key mykey1
```

The following example shows how to configure use of an authentication keychain:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospf)# authentication message-digest keychain mykeychain
```

Related Commands

| Command | Description |
|--|---|
| authentication-key (OSPF), on page 939 | Assigns a password to be used by neighboring routers that are using the simple password authentication of OSPF. |
| message-digest-key, on page 1013 | Specifies a key used with OSPF MD5 authentication. |

authentication-key (OSPF)

To assign a password to be used by neighboring routers that are using the Open Shortest Path First (OSPF) simple password authentication, use the **authentication-key** command in the appropriate mode. To remove a previously assigned OSPF password, use the **no** form of this command.

authentication-key [**clear**|**encrypted**] *password*

no authentication-key

Syntax Description

| | |
|------------------|---|
| clear | (Optional) Specifies that the key be clear text. |
| encrypted | (Optional) Specifies that the key be encrypted using a two-way algorithm. |
| <i>password</i> | Any contiguous string up to 8 characters in length that can be entered from the keyboard. For example, <i>mypswd2</i> . |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the OSPF password parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the OSPF password parameter specified for the process.

If this command is not specified at any level, then no password is specified.

Clear is the default if the **clear** or **encrypted** keyword is not specified.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The password created by this command is inserted directly into the OSPF header when the Cisco IOS XR software originates routing protocol packets. A separate password can be assigned to each network on an individual interface basis. All neighboring routers on the same network must have the same password to be able to exchange OSPF information.

The **authentication-key** command must be used with the **authentication** command. If the **authentication** command is not configured, the password provided by the **authentication-key** command is ignored and no authentication is adopted by the OSPF interface.

**Note**

The **authentication-key** command cannot be used with the **authentication** command when the **message-digest** or **null** keyword is configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure an authentication password as the string yourpass:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# authentication-key yourpass
```

Related Commands

| Command | Description |
|---|--------------------------------|
| authentication (OSPF) , on page 936 | Specifies authentication type. |

auto-cost (OSPF)

To control how the Open Shortest Path First (OSPF) protocol calculates default metrics for the interface, use the **auto-cost** command in the appropriate mode. To revert to the default reference bandwidth, use the **no** form of this command.

auto-cost {reference-bandwidth *mbps*| **disable**}

no auto-cost {reference-bandwidth| **disable**}

Syntax Description

| | |
|--|--|
| reference-bandwidth <i>mbps</i> | Specifies a rate in Mbps (bandwidth). Range is 1 to 4294967. |
| disable | Assigns a cost based on interface type. |

Command Default

mbps : 100 Mbps

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface.

The OSPF metric is calculated as the *mbps* value divided by bandwidth, with *mbps* equal to 108 by default.

If you have multiple links with high bandwidth (such as OC-192), you might want to use a larger number to differentiate the cost on those links. That is, the metric calculated using the default *mbps* value is the same for all high-bandwidth links.

Recommended usage of cost configuration for OSPF interfaces with high bandwidth is to be consistent: Either explicitly configure (by using the **cost** command) or choose the default (by using the **auto-cost** command).

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the reference value for the auto cost calculation to 1000 Mbps:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf)# auto-cost reference-bandwidth 1000
```

Related Commands

| Command | Description |
|---|---|
| cost (OSPF) , on page 954 | Explicitly specifies the cost of the interface (network) for OSPF path calculation. |

capability opaque disable

To prevent Multiprotocol Label Switching traffic engineering (MPLS TE) topology information flooded to the network through opaque LSAs, use the **capability opaque disable** command in the appropriate mode. To restore MPLS TE topology information flooded through opaque LSAs to the network, use the **no** form of the command.

capability opaque disable

no capability opaque disable

Command Default Opaque LSAs are allowed.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **capability opaque disable** command prevents flooded MPLS TE information (Types 1 and 4) through opaque LSAs of all scope (Types 9, 10, and 11).

Control opaque LSA support capability must be enabled for OSPF to support MPLS TE.

The MPLS TE topology information is flooded to the area through opaque LSAs by default.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ospf | read, write |

Examples The following example shows how to prevent OSPF from supporting opaque services:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# capability opaque disable
```

capability vrf-lite (OSPFv2)

To ignore DN bit in LSAs received from peers in the given VRF and to disable automatic ABR status in that VRF, use the **capability vrf-lite** command in OSPFv2 VRF configuration mode. To disable ignoring the DN bit in LSAs and to re-enable automatic ABR status in the VRF, use the **no** form of this command.

capability vrf-lite

no capability vrf-lite

Syntax Description This command has no keywords or arguments.

Command Default Disabled

Command Modes OSPFv2 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **capability vrf-lite** command when CE routers (sometimes called multi-vrf CE routers) are directly connected through interfaces associated with a VRF, but without being connected to other PEs through the MPLS/VPN BGP Backbone. Configure the **capability vrf-lite** command for CE routers.

When OSPFv2 is enabled in a VRF, the router is always an ABR. With the **capability vrf-lite** command is enabled, the router becomes an ABR only if it is connected to area 0 (backbone area), and there are other (non-backbone) areas enabled on this router in the given VRF.



Note Routes may be re-introduced to the VPN backbone when this command is used.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to enable vrf-lite capability for OSPF instance *100* under *vrf1*

```
RP/0/RSP0/CPU0:router#configure  
RP/0/RSP0/CPU0:router(config)#router ospf 100  
RP/0/RSP0/CPU0:router(config-ospf)#vrf vrf1  
RP/0/RSP0/CPU0:router(config-ospf-vrf)#capability vrf-lite
```

clear ospf process

To reset an Open Shortest Path First (OSPF) router process without stopping and restarting it, use the **clear ospf process** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name*| **all**}]] **process**

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF processes are reset. |
| vrf | (Optional) An OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF instance to be reset. |
| all | (Optional) Resets all OSPF VRF instances. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the OSPF router process is reset, OSPF releases all resources allocated, cleans up the internal database, and shuts down and restarts all interfaces that belong to the process.



Note

The **clear ospf process** command may change the router ID unless the OSPF router ID is explicitly configured through the [router-id \(OSPF\)](#), on page 1076 command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to reset all OSPF processes:

```
RP/0/RSP0/CPU0:router# clear ospf process
```

The following example shows how to reset the OSPF 1 process:

```
RP/0/RSP0/CPU0:router# clear ospf 1 process
```

Related Commands

| Command | Description |
|---|--|
| router ospf , on page 1078 | Configures an OSPF routing process. |
| router-id (OSPF) , on page 1076 | Configures a router ID for the OSPF process. |

clear ospf redistribution

To clear all routes redistributed from other protocols out of the Open Shortest Path First (OSPF) routing table, use the **clear ospf redistribution** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name*| **all**}]] **redistribution**

Syntax Description

| | |
|---------------------|--|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared. |
| vrf | (Optional) OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF instance to be reset. |
| all | (Optional) Resets all OSPF VRF instances. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospf redistribution** command to cause the routing table to be read again. OSPF regenerates and sends Type 5 and Type 7 link-state advertisements (LSAs) to its neighbors. If an unexpected route has appeared in the OSPF redistribution, using this command corrects the issue.



Note

Use of this command can cause a significant number of LSAs to flood the network. We recommend that you use this command with caution.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to clear all redistributed routes across all processes from other protocols:

```
RP/0/RSP0/CPU0:router# clear ospf redistribution
```

clear ospf routes

To clear all Open Shortest Path First (OSPF) routes from the OSPF routing table, use the **clear ospf routes** command in EXEC configuration mode.

clear ospf [*process-name* [**vrf** {*vrf-name*| **all**}]] **routes**

Syntax Description

| | |
|---------------------|--|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF routes are cleared. |
| vrf | (Optional) OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF instance to be reset. |
| all | (Optional) Resets all OSPF VRF instances. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to clear all OSPF routes from the OSPF routing table and recompute valid routes. When the OSPF routing table is cleared, OSPF routes in the global routing table are also recalculated.

```
RP/0/RSP0/CPU0:router# clear ospf routes
```

Related Commands

| Command | Description |
|--|-------------------------------------|
| router ospf , on page 1078 | Configures an OSPF routing process. |

clear ospf statistics

To clear the Open Shortest Path First (OSPF) statistics of neighbor state transitions, use the **clear ospf statistics** command in EXEC configuration mode.

```
clear ospf [process-name [vrf {vrf-name| all}]] statistics [neighbor [type interface-path-id] [ip-address ]]
```

Syntax Description

| | |
|--------------------------|--|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPF statistics of neighbor state transitions are cleared. |
| vrf | (Optional) OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF instance to be reset. |
| all | (Optional) Resets all OSPF VRF instances. |
| neighbor | (Optional) Clears the state transition counters of the specified neighbor only. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (Optional) Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>ip-address</i> | (Optional) IP address of a specified neighbor for whom you want to clear the state transition counter. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospf statistics** command to reset OSPF counters. Reset is useful to detect changes in counter values.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to reset the OSPF transition state counters for all neighbors on Packet-over-SONET/SDH (POS) interface 0/2/0/0:

```
RP/0/RSP0/CPU0:router# clear ospf statistics neighbor POS 0/2/0/0
```

Related Commands

| Command | Description |
|--|-------------------------------------|
| router ospf , on page 1078 | Configures an OSPF routing process. |

cost (OSPF)

To explicitly specify the interface (network) for Open Shortest Path First (OSPF) path calculation, use the **cost** command in the appropriate mode. To remove the cost, use the **no** form of this command.

cost *cost*

no cost

Syntax Description

| | |
|-------------|---|
| <i>cost</i> | Unsigned integer value expressed as the link-state metric. Range is 1 to 65535. |
|-------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is calculated by the **auto-cost** command.

Command Modes

Interface configuration

Area configuration

Router configuration

VRF configuration

Multi-area configuration

Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The link-state metric is advertised as the link cost in the router link advertisement. Cisco IOS XR software does not support type of service (ToS), so you can assign only one cost for each interface.

In general, the path cost is calculated using the following formula:

$108 / \text{bandwidth}$ (the default auto cost is set to 100 Mbps)

This calculation is the default reference bandwidth used by the auto-costing calculation which establishes the interface auto-cost. The **auto-cost** command can set this reference bandwidth to some other value. The **cost** command is used to override the auto-costing calculated default value for interfaces.

Using this formula, the default path cost is 1 for any interface that has a link bandwidth of 100 Mbps or higher. If this value does not suit the network, configure the reference bandwidth for auto calculating costs based on the link bandwidth.

The value set by the **cost** command overrides the cost resulting from the **auto-cost (OSPF)** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the cost value to 65 for GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# cost 65
```

Related Commands

| Command | Description |
|---|--|
| auto-cost (OSPF), on page 941 | Controls how the OSPF protocol calculates default metrics for the interface. |

cost-fallback (OSPF)

To apply higher cost than the normal interface cost when the cumulative bandwidth of a bundle interface goes below the threshold specified and to revert to the original cost if the cumulative bandwidth goes above the configured threshold, use the **cost-fallback** command. To remove the cost-fallback, use the **no** form of this command.

cost-fallback cost threshold bandwidth

no cost-fallback

Syntax Description

| | |
|------------------------------|--|
| <i>cost</i> threshold | Unsigned integer value expressed as the link-state metric. Range is 1 to 65535, but typically, cost-fallback value is supposed to be set to a value higher than the normal cost. |
| <i>bandwidth</i> | Unsigned integer value expressed in Mbits per second. Range is 1 to 4294967. |

Command Default

If this command is not specified in interface configuration mode, the currently effective interface cost takes effect even when the cumulative bandwidth goes down below the maximum bandwidth. Unlike the interface cost command, this cost-fallback command is available only under interface configuration mode; it is not available in area or process level. Unlike other interface specific parameters, no inheritance will take place from area or process level if this command is not specified at interface level.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The fallback cost must be set to a higher value than the normal interface cost. The motivation of setting the fallback cost is to cost out an interface or disfavor an interface without shutting it down when its cumulative bandwidth goes below the user specified threshold, so that the traffic can take an alternative path. The normal interface cost will take over when the cumulative bandwidth reaches or exceeds user-specified threshold.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the cost-fallback value:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router (config) #router ospf 100
RP/0/RSP0/CPU0:router (config-ospf) #router-id 2.2.2.2
RP/0/RSP0/CPU0:router (config-ospf) #area 0
RP/0/RSP0/CPU0:router (config-ospf-ar) #interface bundle-a
RP/0/RSP0/CPU0:router (config-ospf-ar-if) #cost-fallback 1000 threshold 300
```

Related Commands

| Command | Description |
|--|--|
| auto-cost (OSPF) , on page 941 | Controls how the OSPF protocol calculates default metrics for the interface. |
| cost (OSPF) , on page 954 | Specifies the cost of the interface (network) for OSPF path calculation. |

database-filter all out (OSPF)

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) interface, use the **database-filter all out** command in the appropriate mode. To restore the forwarding of LSAs to the interface, use the **disable** form of the command.

database-filter all out [**disable**| **enable**]

Syntax Description

| | |
|----------------|--------------------------------|
| disable | (Optional) Disables filtering. |
| enable | (Optional) Enables filtering. |

Command Default

The database filter is disabled.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration
 Multi-area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **database-file all out** command to perform the same function that the [neighbor database-filter all out](#), on page 1033 command performs on a neighbor basis.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to prevent flooding of OSPF LSAs to broadcast, nonbroadcast, and point-to-point networks reachable through GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# database-filter all out
```

Related Commands

| Command | Description |
|--|---|
| neighbor database-filter all out, on page 1033 | Filters outgoing LSAs to an OSPF interface. |

dead-interval (OSPF)

To set the interval after which a neighbor is declared dead when no hello packets are observed, use the **dead-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

dead-interval *seconds*

no dead-interval

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Integer that specifies the interval (in seconds). Range is 1 to 65535. The value must be the same for all nodes on the network. |
|----------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified at any level, then the dead interval is four times the interval set by the **hello-interval** (OSPF) command.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The dead interval value must be the same for all routers and access servers on a specific network.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the OSPF dead interval to 40 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# dead-interval 40
```

Related Commands

| Command | Description |
|--|---|
| hello-interval (OSPF), on page 993 | Specifies the interval between hello packets that the Cisco IOS XR software sends on the interface. |

default-cost (OSPF)

To specify a cost for the default summary route sent into a stub area or not-so-stubby area (NSSA), use the **default-cost** command in area configuration mode. To remove the assigned default route cost, use the **no** form of this command.

default-cost *cost*

no default-cost *cost*

Syntax Description

| | |
|-------------|---|
| <i>cost</i> | Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number. |
|-------------|---|

Command Default

cost : 1

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub** command in the area submode. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to assign a default cost of 20 to a stub area. The GigabitEthernet interface 0/4/0/3 is also configured in the stub area:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 201  
RP/0/RSP0/CPU0:router(config-ospf)# area 10.15.0.0  
RP/0/RSP0/CPU0:router(config-ospf-ar)# stub  
RP/0/RSP0/CPU0:router(config-ospf-ar)# default-cost 20  
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/4/0/3
```

Related Commands

| Command | Description |
|---|---------------------------------|
| stub (OSPF), on page 1153 | Defines an area as a stub area. |

default-information originate (OSPF)

To generate a default external route into an Open Shortest Path First (OSPF) routing domain, use the **default-information originate** command in the appropriate mode. To disable this feature, use the **no** form of this command.

default-information originate [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**route-policy** *policy-name*] [**tag** *tag-value*]

no default-information originate

Syntax Description

| | |
|--|---|
| always | (Optional) Always advertises the default route regardless of whether the routing table has a default route. |
| metric <i>metric-value</i> | (Optional) Specifies the metric used for generating the default route. The default metric value is 1. Range is 1 to 16777214. |
| metric-type <i>type-value</i> | (Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: 1—Type 1 external route 2—Type 2 external route |
| tag <i>tag-value</i> | (Optional) 32-bit dotted-decimal value attached to each external route. This is not used by the OSPF protocol itself. It may be used to communicate information between autonomous system boundary routers (ASBRs). If a tag is not specified, then zero (0) is used. |
| route-policy <i>policy-name</i> | (Optional) Specifies that a routing policy be used and the routing policy name. |

Command Default

When you do not use this command in router configuration mode, no default external route is generated into an OSPF routing domain.

metric-value : 1

type-value : 2

tag-value: 0

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Whenever you use the **redistribute** or **default-information originate** command to redistribute routes into an OSPF routing domain, the software automatically becomes an Autonomous System Boundary Router (ASBR). However, an ASBR does not, by default, generate a default route into the OSPF routing domain. The software still must have a default route for itself before it generates one, except when you have specified the **always** keyword.

The **default-information originate** route-policy attach point conditionally injects the default route 0.0.0.0/0 into the OSPF link-state database, and is done by evaluating the attached policy. If any routes specified in the policy exist in the global RIB, then the default route is inserted into the link-state database. If there is no match condition specified in the policy, the policy passes and the default route is generated into the link-state database.

For information about the default-information originate attach point, see the *OSPF Policy Attach Points* section in the *Implementing Routing Policy* chapter in *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify a metric of 100 for the default route redistributed into the OSPF routing domain and an external metric type of Type 1:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)#redistribute igmp 108 metric 100
RP/0/RSP0/CPU0:router(config-ospf)#default-information originate metric 100 metric-type 1
```

Related Commands

| Command | Description |
|--|---|
| redistribute (OSPF) , on page 1067 | Redistributes routes from one routing domain into a specified OSPF process. |

default-metric (OSPF)

To set default metric values for routes redistributed from another protocol into the Open Shortest Path First (OSPF) protocol, use the **default-metric** command in the appropriate mode. To return to the default state, use the **no** form of this command.

default-metric *value*

no default-metric *value*

Syntax Description

| | |
|--------------|--|
| <i>value</i> | Default metric value appropriate for the specified routing protocol. Range is 1 to 16777214. |
|--------------|--|

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause the current routing protocol to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use a default metric to provide a reasonable substitute and enable the redistribution to proceed.

The default-metric value configured in OSPF configuration does not apply to connected routes that are redistributed to OSPF using the **redistribute connected** command. To set a non-default metric for connected routes, configure OSPF with the **redistribute connected metric** *metric-value* command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to advertise Intermediate System-to-Intermediate System (IS-IS) protocol-derived routes into OSPF and assign a metric of 10:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# default-metric 10
RP/0/RSP0/CPU0:router(config-ospf)# redistribute isis IS-IS_osp
```

Related Commands

| Command | Description |
|---|---|
| redistribute (OSPF), on page 1067 | Redistributes routes from one routing domain into a specified OSPF process. |

demand-circuit (OSPF)

To configure the Open Shortest Path First (OSPF) protocol to treat the interface as an OSPF demand circuit, use the **demand-circuit** command in the appropriate mode. To remove the demand circuit designation from the interface, use the **no** form of this command.

demand-circuit [**disable**| **enable**]

no demand-circuit

Syntax Description

| | |
|----------------|--|
| disable | (Optional) Disables the interface as an OSPF demand circuit. |
| enable | (Optional) Enables the interface as an OSPF demand circuit. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the circuit is not a demand circuit.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On point-to-point interfaces, only one end of the demand circuit must be configured with this command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. Use the **demand-circuit** command to allow the underlying data link layer to be closed when the topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows how to set the configuration for an OSPF demand circuit:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf)# demand-circuit
```

disable-dn-bit-check

To specify that down bits should be ignored, use the **disable-dn-bit-check** command in VPN routing and forwarding (VRF) configuration mode. To specify that down bits should be considered, use the **no** form of this command.

disable-dn-bit-check

no disable-dn-bit-check

Command Default Down bits are considered.

Command Modes VRF configuration mode

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to specify that down bits be ignored:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# vrf v1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# disable-dn-bit-check
```

distance (OSPF)

To define an administrative distance, use the **distance** command in an appropriate configuration mode. To remove the **distance** command from the configuration file and restore the system to its default condition in which the software removes a distance definition, use the **no** form of this command.

distance *weight* [*ip-address wildcard-mask* [*access-list-name*]]

no distance *weight ip-address wildcard-mask* [*access-list-name*]

Syntax Description

| | |
|-------------------------|---|
| <i>weight</i> | Administrative distance. Range is 10 to 255. Used alone, the <i>weight</i> argument specifies a default administrative distance that the software uses when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. lists the default administrative distances. Table 88: Default Administrative Distances, on page 972 |
| <i>ip-address</i> | (Optional) IP address in four-part, dotted-decimal notation. |
| <i>wildcard-mask</i> | (Optional) Wildcard mask in four-part, dotted decimal format. A bit set to 1 in the <i>mask</i> argument instructs the software to ignore the corresponding bit in the address value. |
| <i>access-list-name</i> | (Optional) Name of an IP access list to be applied to incoming routing updates. |

Command Default

If this command is not specified, then the administrative distance is the default, as specified in [Table 88: Default Administrative Distances, on page 972](#).

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An administrative distance is an integer from 10 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. Weight values are subjective; no quantitative method exists for choosing weight values.

If an access list is used with this command, it is applied when a network is being inserted into the routing table. This behavior allows you to filter networks based on the IP prefix supplying the routing information. For example, you could filter possibly incorrect routing information from networking devices not under your administrative control.

The order in which you enter **distance** commands can affect the assigned administrative distances in unexpected ways (see the “Examples” section for further clarification).

This table lists default administrative distances.

Table 88: Default Administrative Distances

| Route Source | Default Distance |
|-------------------------------|------------------|
| Connected interface | 0 |
| Static route out on interface | 0 |
| State route to next-hop | 1 |
| EIGRP Summary Route | 5 |
| External BGP | 20 |
| Internal EIGRP | 90 |
| OSPF | 110 |
| IS-IS | 115 |
| RIP version 1 and 2 | 120 |
| External EIGRP | 170 |
| Internal BGP | 200 |
| Unknown | 255 |

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

In the following example, the **router ospf** command sets up OSPF routing instance1. The first **distance** command sets the default administrative distance to 255, which instructs the software to ignore all routing updates from networking devices for which an explicit distance has not been set. The second **distance**

command sets the administrative distance for all networking devices on the Class C network 192.168.40.0 0.0.0.255 to 90.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# distance 255
RP/0/RSP0/CPU0:router(config-ospf)# distance 90 192.168.40.0 0.0.0.255
```

Related Commands

| Command | Description |
|--|--|
| distance bgp | Allows the use of external, internal, and local administrative distances that could be a better route to a BGP node. |
| distance ospf | Allows the use of external, internal, and local administrative distances that could be a better route to an OSPF node. |
| router ospf , on page 1078 | Configures the OSPF routing process. |

distance ospf

To define Open Shortest Path First (OSPF) route administrative distances based on route type, use the **distance ospf** command in router configuration mode. To restore the default value, use the **no** form of this command.

distance ospf {**intra-area**|**inter-area**|**external**} *distance*

no distance ospf

Syntax Description

intra-area | **inter-area** | **external**

Sets the type of area. It can be one of the following values:

intra-area —All routes within an area.

inter-area —All routes from one area to another area.

external —All routes from other routing domains, learned by redistribution.

Any combination of the above areas is allowed.

distance

Route administrative distance.

Command Default

distance : 110

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must specify one of the keywords.

Use the **distance ospf** command to perform the same function as the **distance** command used with an access list. However, the **distance ospf** command sets a distance for an entire group of routes, rather than a specific route that passes an access list.

A common reason to use the **distance ospf** command is when you have multiple OSPF processes with mutual redistribution, and you want to prefer internal routes from one over external routes from the other.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the external distance to 200, making the route less reliable:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# distance ospf external 200
RP/0/RSP0/CPU0:router(config-ospf)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# distance ospf external 200
```

Related Commands

| Command | Description |
|--|-------------------------------------|
| disable-dn-bit-check , on page 970 | Defines an administrative distance. |

distribute-list

To filter networks received or transmitted in Open Shortest Path First (OSPF) updates, use the **distribute-list** command in the appropriate mode. To change or cancel the filter, use the **no** form of this command.

distribute-list {*access-list-name* {**in**|**out** [**bgp number**|**connected**|**ospf instance**|**static**]}}| **route-policy** *route-policy-name in*}

no distribute-list {*access-list-name* {**in**|**out**}}| **route-policy** *route-policy-name in*}

Syntax Description

| | |
|---|---|
| <i>access-list-name</i> | Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates. |
| in | Applies the access list or route-policy to incoming routing updates. |
| out | Applies the access list to outgoing routing updates. The out keyword is available only in router configuration mode. |
| bgp | (Optional) Applies the access list to BGP routes. |
| connected | (Optional) Applies the access list to connected routes. |
| ospf | (Optional) Applies the access list to OSPF routes (not the current OSPF process). |
| static | (Optional) Applies the access list to statically configured routes. |
| route-policy <i>route-policy-name</i> | Specifies the route-policy to filter OSPF prefixes. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the distribute list parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the distribute list parameter specified for the process.

If this command is not specified at any level, then the distribute list is disabled.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration
 Multi-area configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.1 | The route-policy <i>route-policy-name</i> keyword and argument were added to allow use of route policies to filter OSPF prefixes. |
| Release 4.3.1 | Support was added for "if tag..." statements in distribute-list in <i>route-policy</i> . |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **distribute-list** command to limit which OSPF routes are installed on this router. The **distribute-list** command does not affect the OSPF protocol itself.

The **distribute-list in** is configurable at instance (process), area, and interface levels. Regular OSPF configuration inheritance applies. Configuration is inherited from instance > area > interface levels.

Use the **route-policy** *route-policy-name* keyword and argument to allow use of route policies to filter OSPF prefixes.

**Note**

Either an access-list, or a route-policy can be used in a single command, not both. Configuring the command with access-list removes the route-policy configuration, and vice versa.

The "if tag..." statements can be used in **distribute-list in** *route-policy*. The matching on route tag supports operators "eq/ge/is/le". Operator "in" is not supported.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to prevent OSPF routes from the 172.17.10.0 network from being installed if they are learned in area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# ipv4 access-list 3
RP/0/RSP0/CPU0:router(config-ipv4-acl)# deny 172.17.10.0 0.0.0.255
RP/0/RSP0/CPU0:router(config-ipv4-acl)# permit any any
!
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# distribute-list 3 in
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
```

domain-id (OSPF)

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove an OSPF VRF domain ID, use the **no** form of this command.

domain-id [secondary] type [0005| 0105| 0205| 8005] value *value*

no domain-id [secondary] type [0005| 0105| 0205| 8005] value *value*

Syntax Description

| | |
|---------------------------|--|
| secondary | (Optional) OSPF secondary domain ID. |
| type | Primary OSPF domain ID in hex format. |
| value <i>value</i> | OSPF domain ID value in hex format (six octets). |

Command Default

No domain ID is specified.

Command Modes

VRF configuration mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An OSPF domain id must be explicitly configured. The OSPF domain ID helps OSPF determine how to translate a prefix received through Border Gateway Protocol (BGP) from the remote provider edge (PE). If the domain IDs match, OSPF generates a Type 3 link state advertisement (LSA). If the domain IDs do not match, OSPF generates a Type 5 LSA.

There is only one primary domain ID. There can be multiple secondary domain IDs.



Note

When an IOS XR router and an IOS router are configured as peers, the two Domain IDs must match. Manually configure the IOS XR Domain ID value to match the IOS default Domain ID value. This ensures that the routes have route code "OIA" because they are learned as inter-area routes. If the Domain IDs do not match, the routes have route code, "O-E2" because they are learned as external routes. Use the **show ip ospf** command to get the OSPF Domain ID from the IOS router. Then, set the IOS XR Domain ID to the same value using the **domain-id** command.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows how to specify a domain ID:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf o1  
RP/0/RSP0/CPU0:router(config-ospf)# vrf v1  
RP/0/RSP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABCCDDEEFF
```

domain-tag

To specify the Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) domain tag, use the **domain-tag** command in VRF configuration mode. To remove an OSPF VRF domain tag, use the **no** form of this command.

domain-tag *tag*

no domain-tag

Syntax Description

| | |
|------------|--|
| <i>tag</i> | OSPF domain tag as a 32-bit value. The valid range is 0 to 4294967295. |
|------------|--|

Command Default

No OSPF VRF domain tag is specified.

Command Modes

VRF configuration mode

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The domain tag is added to any Type 5 link state advertisements (LSAs) generated as a result of VPN-IP routes received from Border Gateway Protocol (BGP). The domain-tag is derived from BGP autonomous system number (ASN).

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify the domain tag:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 01
RP/0/RSP0/CPU0:router(config-ospf)# domain-tag 234
```

fast-reroute (OSPFv2)

To enable IP fast reroute loop-free alternate (LFA) computation, use the **fast-reroute** command in the appropriate OSPF configuration mode. To disable the IP fast reroute loop-free alternate computation, use the **no** form of this command.

To disable loop-free alternate computation that is enabled on a higher level, use the **fast-reroute** command with **disable** keyword.

fast-reroute {per-link| per-prefix} [disable]

no fast-reroute

Syntax Description

| | |
|-------------------|---|
| per-link | Enables per-link loop-free alternate computation. |
| per-prefix | Enables per-prefix loop-free alternate computation. |
| disable | (Optional) Disables loop-free alternate computation that was enabled on a higher level. |

Command Default

IP fast-reroute LFA computation is disabled.

Command Modes

Area configuration
Interface configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Only one mode of computation can be configured on an interface - per-link or per-prefix. Different modes of computations can be enabled on different interfaces; one set of interface using per-link and other set using per-prefix computation. Based on the outgoing interface of the primary path, per-link or per-prefix backup path will be computed.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | ospf | read, write |

Examples

This example shows how to enable per-link computation of loop-free alternates under interface POS 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface POS 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link
```

This example shows how to enable per-prefix computation of loop-free alternates under area 0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
```

This example shows how to disable computation of loop-free alternates that was configured under area 0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)#fast-reroute per-prefix
RP/0/RSP0/CPU0:router(config-ospf-ar)#interface POS 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)#fast-reroute disable
```

fast-reroute per-link exclude interface

To excludes specified interface to be used as a backup during (IPFRR) loop-free alternate (LFA) computation, use the **fast-reroute per-link exclude interface** command, in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-link exclude interface *type interface-path-id*

no fast-reroute per-link exclude interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No interfaces are excluded.

Command Modes

Interface configuration
Area configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to exclude an interface from IP fast reroute loop-free alternate (LFA) computation:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# fast-reroute per-link exclude interface POS 0/3/2/1
```

Related Commands

| Command | Description |
|--|--|
| fast-reroute (OSPFv2), on page 981 | Enables IP fast reroute loop-free alternate (LFA) computation. |

fast-reroute per-prefix exclude interface (OSPFv2)

To exclude interface to be used as a backup path from fast-reroute loop-free alternate per-prefix computation, use the **fast-reroute per-prefix exclude interface** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix exclude interface *type interface-path-id*

no fast-reroute per-prefix exclude interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|--|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No interfaces are excluded.

Command Modes

Interface configuration
Area configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Backup paths via the excluded interfaces will not be computed.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to exclude interface POS0/6/0/1 from being used as a backup path:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix exclude interface POS0/6/0/1
```

fast-reroute per-prefix lfa-candidate (OSPFv2)

To add interfaces to the LFA candidate list, use the **fast-reroute per-prefix lfa-candidate** command in interface configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix lfa-candidate [*interface-name*]

no fast-reroute per-prefix lfa-candidate [*interface-name*]

Syntax Description

| | |
|-----------------------|---|
| <i>interface-name</i> | Specifies name of the interface to add to the LFA candidate list. |
|-----------------------|---|

Command Default

No interfaces are added to the candidate list.

Command Modes

Interface configuration
Area configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to add an interface to LFA candidates:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix lfa-candidate interface POS0/6/0/0
```

fast-reroute per-prefix remote-lfa (OSPFv2)

To configure fast-reroute per-prefix remote loop-free alternate (rLFA) computation for an OSPFv2 process, use the **fast-reroute per-prefix remote-lfa** command in the appropriate OSPF configuration mode. To disable this feature, use the **no** form of this command.

```
fast-reroute per-prefix remote-lfa {disable| maximum-cost path-cost| tunnel mpls-ldp}
no fast-reroute per-prefix remote-lfa
```

Syntax Description

| | |
|--------------------------------------|---|
| remote-lfa | Enables remote LFA backup computation |
| maximum-cost <i>path-cost</i> | Sets the cost option to limit the range of remote LFAs. Range for path-cost is 1 to 4294967295. |
| tunnel mpls-ldp | Enables remote LFA computation using tunnel interfaces. |
| disable | Selectively disables remote LFA calculation under one or more areas. |

Command Default

Remote LFA FRR computation is disabled.

Command Modes

Router configuration
Area configuration
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Fast-reroute per-prefix LFA must be enabled for remote LFA to be operational. The OSPF configuration hierarchy at process level and area level are applicable for remote LFA configuration. For example, it is possible to enable remote-LFA for all OSPF areas and selectively disable (using **disable** keyword) the computation under one or more area.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to enable fast-reroute per-prefix remote LFA computation for tunnel interfaces:

```
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix remote-lfa tunnel mpls-ldp
```

This example shows how to configure cost option to limit the range of remote LFAS:

```
RP/0/RSP0/CPU0:router(config)#router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix remote-lfa maximum-cost 2
```

Related Commands

| Command | Description |
|--|--|
| fast-reroute (OSPFv2), on page 981 | Enables IP fast reroute loop-free alternate (LFA) computation. |

fast-reroute per-prefix use-candidate-only (OSPFv2)

To restrict the backup interfaces to those that are present on the LFA candidate list, use the **fast-reroute per-prefix use-candidate-only** command in router OSPF configuration mode. To disable this feature, use the **no** form of this command.

fast-reroute per-prefix use-candidate-only [enable| disable]

fast-reroute per-prefix use-candidate-only

Syntax Description

| | |
|----------------|---|
| enable | Enables backup selection from candidate-list only. |
| disable | Disables backup selection from candidate-list only. |

Command Default

Disabled.

Command Modes

Router OSPF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to restrict the backup interfaces to those that are present on the LFA candidate list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#fast-reroute per-prefix use-candidate-only
```

flood-reduction (OSPF)

To suppress the unnecessary flooding of link-state advertisements (LSAs) in stable topologies, use the **flood-reduction** command in the appropriate mode. To remove this functionality from the configuration, use the **no** form of this command.

flood-reduction [**enable**| **disable**]

no flood-reduction [**enable**| **disable**]

Syntax Description

| | |
|----------------|--|
| enable | (Optional) Turns on this functionality at a specific level. |
| disable | (Optional) Turns off this functionality at a specific level. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All routers supporting the OSPF demand circuit are compatible and can interact with routers supporting flooding reduction.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# flood-reduction
```

Related Commands

| Command | Description |
|---|--|
| show ospf interface, on page 1106 | Displays OSPF-related interface information. |
| show ospf neighbor, on page 1117 | Displays OSPF neighbor information on an individual interface basis. |

hello-interval (OSPF)

To specify the interval between consecutive hello packets that are sent on the Open Shortest Path First (OSPF) interface, use the **hello-interval** command in the appropriate mode. To return to the default time, use the **no** form of this command.

hello-interval *seconds*

no hello-interval

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Interval (in seconds). The value must be the same for all nodes on a specific network. Range is 1 to 65535. |
|----------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (nonbroadcast).

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic occurs. This value must be the same for all routers and access servers on a specific network.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the interval between hello packets to 15 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# hello-interval 15
```

Related Commands

| Command | Description |
|--|--|
| dead-interval (OSPF) , on page 960 | Sets the time period for which hello packets are suspended before neighbors declare the router down. |

ignore lsa mospf

To suppress the sending of syslog messages when the router receives link-state advertisement (LSA) Type 6 multicast Open Shortest Path First (MOSPF) packets, which are unsupported, use the **ignore lsa mospf** command in an appropriate configuration mode. To restore the sending of syslog messages, use the **no** form of this command.

ignore lsa mospf

no ignore lsa mospf

This command has no keywords or arguments.

Command Default

When you do not specify this command in router configuration mode, each MOSPF packet received by the router causes the router to send a syslog message.

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Cisco routers do not support LSA Type 6 (MOSPF), and they generate syslog messages if they receive such packets. If the router is receiving many MOSPF packets, you might want to configure the router to ignore the packets and thus prevent a large number of syslog messages.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the router to suppress the sending of syslog messages when it receives MOSPF packets:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# ignore lsa mospf
```

interface (OSPF)

To define the interfaces on which the Open Shortest Path First (OSPF) protocol runs, use the **interface** command in area configuration mode. To disable OSPF routing for interfaces, use the **interface** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

When you do not specify this command in configuration mode, OSPF routing for interfaces is not enabled.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an area. The interface remains associated with the area even when the IP address of the interface changes.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how the OSPF routing process 109 defines four OSPF areas (0, 2, 3, and 10.9.50.0), and associates an interface with each area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router (config)# router ospf 109
RP/0/RSP0/CPU0:router (config-ospf)# area 0
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 4/0/0/3
!
RP/0/RSP0/CPU0:router (config-ospf)# area 2
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 0/1/0/3
!
RP/0/RSP0/CPU0:router (config-ospf)# area 3
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 3/0/0/2
!
RP/0/RSP0/CPU0:router (config-ospf)# area 10.9.50.0
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet 3/0/0/1
```

log adjacency changes (OSPF)

To configure the router to send a syslog message when the state of an Open Shortest Path First (OSPF) neighbor changes, use the **log adjacency changes** command in router configuration mode. To turn off this function, use the **disable** keyword. To log all state changes, use the **detail** keyword.

log adjacency changes {**detail**|**disable**}

Syntax Description

| | |
|----------------|--|
| detail | Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes. |
| disable | Disables sending adjacency change messages. |

Command Default

The router sends a syslog message when the state of an OSPF neighbor changes.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **log adjacency changes** command to display high-level changes to the state of the peer relationship. Configure this command if you want to know about OSPF neighbor changes.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the software to send a syslog message for any OSPF neighbor state changes:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# log adjacency changes detail
```

loopback stub-network

To enable advertising loopback as stub networks, use the **loopback stub-network** command in an appropriate configuration mode. To disable advertising loopback as stubnetworks, use the **no** form of this command.

loopback stub-network [**enable**| **disable**]

no loopback stub-network

Syntax Description

| | |
|----------------|---|
| enable | (Optional) Enables advertising loopbacks as stub networks. |
| disable | (Optional) Disables advertising loopbacks as stub networks. |

Command Default

By default, OSPF advertises loopbacks as stub hosts.

Command Modes

OSPF interface configuration
 OSPF router configuration
 OSPF area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In the interface submode, the command can be enabled only on loopback interfaces.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable advertising loopback as a stub network, under OSPF interface configuration:

```
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#loopback stub-network enable
```

Related Commands

| Command | Description |
|--|---|
| show ospf interface , on page 1106 | Displays Open Shortest Path First (OSPF) interface information. |

max-lsa

To limit the number of nonself-generated link-state advertisements (LSAs) that an Open Shortest Path First (OSPF) routing process can keep in the OSPF link-state database (LSDB), use the **max-lsa** command in router configuration mode. To remove the limit of non self-generated LSAs that an OSPF routing process can keep in the OSPF LSDB, use the **no** form of this command.

max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]

no max-lsa *max* [*threshold*] [**warning-only**] [**ignore-time** *value*] [**ignore-count** *value*] [**reset-time** *value*]

Syntax Description

| | |
|----------------------------------|--|
| <i>max</i> | Maximum number of nonself-generated LSAs the OSPF process can keep in the OSPF LSDB. |
| <i>threshold</i> | (Optional) The percentage of the maximum LSA number, as specified by the maximum-number argument, at which a warning message is logged. The default is 75 percent. |
| warning-only | (Optional) Specifies that only a warning message is sent when the maximum limit for LSAs is exceeded. Disabled by default. |
| ignore-time <i>value</i> | (Optional) Specifies the time, in minutes, to ignore all neighbors after the maximum limit of LSAs has been exceeded. The default is 5 minutes. |
| ignore-count <i>value</i> | (Optional) Specifies the number of times the OSPF process can consecutively be placed into the ignore state. The default is 5 times. |
| reset-time <i>value</i> | (Optional) Specifies the time, in minutes, after which the ignore count is reset to zero. The default is 2 times ignore-time . |

Command Default

Disabled

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This command allows you to protect the OSPF routing process from the large number of received LSAs that can result from a misconfiguration on another router in the OSPF domain (for example, the redistribution of a large number of IP prefixes to OSPF).

When this feature is enabled, the router keeps count of the number of all received (nonself-generated) LSAs. When the configured *threshold* value is reached, an error message is logged. When the configured *max* number of received LSAs is exceeded, the router stops accepting new LSAs.

If the count of received LSAs is higher than the configured *max* number after one minute, the OSPF process disables all adjacencies in the given context and clears the OSPF database. This state is called the ignore state. In this state, all OSPF packets received on all interfaces belonging to the OSPF instance are ignored and no OSPF packets are generated on its interfaces. The OSPF process remains in the ignore state for the duration of the configured **ignore-time**. When the **ignore-time** expires, the OSPF process returns to normal operation and starts building adjacencies on all its interfaces.

To prevent the OSPF instance from endlessly oscillating between its normal state and the ignore state, as a result of the LSA count immediately exceeding the *max* number again after it returns from the ignore state, the OSPF instance keeps a count of how many times it has been in the ignore state. This counter is called the **ignore-count**. If the **ignore-count** exceeds its configured value, the OSPF instance remains in the ignore state permanently.

To return the OSPF instance to its normal state, you must issue the **clear ip ospf** command. The **ignore-count** is reset to zero if the LSA count does not exceed the *max* number again during the time configured by the **reset-time** keyword.

If you use the **warning-only** keyword, the OSPF instance never enters the ignore state. When LSA count exceeds the *max* number, the OSPF process logs an error message and the OSPF instance continues in its normal state operation.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the OSPF instance to accept 12000 nonself-generated LSAs in the global routing table, and 1000 nonself-generated LSAs in VRF V1.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 0
RP/0/RSP0/CPU0:router(config-ospf)# max-lsa 12000
RP/0/RSP0/CPU0:router(config-ospf)# vrf V1
RP/0/RSP0/CPU0:router(config-ospf)# max-lsa 1000
```

The following example shows how to display the current status of the OSPF instance:

```
RP/0/RSP0/CPU0:router# show ospf 0

Routing Process "ospf 0" with ID 10.0.0.2
NSR (Non-stop routing) is Disabled
Supports only single TOS(TOS0) routes
Supports opaque LSA
It is an area border router
Maximum number of non self-generated LSA allowed 12000
Current number of non self-generated LSA 1
Threshold for warning message 75%
```

```
Ignore-time 5 minutes, reset-time 10 minutes  
Ignore-count allowed 5, current ignore-count 0
```

Related Commands

| Command | Description |
|---|---|
| show ospf, on page 1084 | Displays general information about Open Shortest Path First (OSPF) routing processes. |

max-metric

To configure the Open Shortest Path First (OSPF) protocol to signal other networking devices not to prefer the local router as an intermediate hop in their shortest path first (SPF) calculations, use the **max-metric** command in router configuration mode. To disable this function, use the **no** form of this command.

max-metric router-lsa [*external-lsa overriding metric*] [**include-stub**] [**on-proc-migration**] [**on-proc-restart**] [**on-startup**] [**on-switchover**] [**wait-for-bgp**] [**summary-lsa**]

no max-metric router-lsa

Syntax Description

| | |
|--|---|
| router-lsa | Always originates router link-state advertisements (LSAs) with the maximum metric. |
| external-lsa <i>overriding metric</i> | (Optional) Overrides the external-lsa metric with the max-metric value. The <i>overriding metric</i> argument specifies the number of in-summary-LSAs. The range is 1 to 16777215. The default is 16711680. |
| include-stub | (Optional) Advertises stub links in router-LSA with the max-metric value (0xFFFF). |
| on-proc-migration <i>time</i> | (Optional) Sets the maximum metric temporarily after a process migration to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-proc-restart <i>time</i> | (Optional) Sets the maximum metric temporarily after a process restart to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-startup <i>time</i> | (Optional) Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. |
| on-switchover <i>time</i> | (Optional) Sets the maximum metric temporarily after a switchover to originate router-LSAs with the max-metric value. The <i>time</i> range is 5 to 86400 seconds. Note OSPF will not populate maximum metric on the router's generated LSAs, when the OSPF routing process is configured to support Nonstop Routing (NSR) or Nonstop Forwarding/Graceful restart (NSF/GR). |
| wait-for-bgp | (Optional) Causes OSPF to originate router LSAs with the maximum metric and allows Border Gateway Protocol (BGP) to decide when to start originating router LSAs with a normal metric instead of the maximum metric. |
| summary-lsa | (Optional) specifies the number of in summary-LSAs. The range is 1 to 16777215. The default is 16711680. |

Command Default

Router LSAs are originated with normal link metrics.

overriding-metric :16711680

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **max-metric** command to cause the software to originate router LSAs with router link metrics set to LSInfinity (0XFFFF). This feature can be useful in Internet backbone routers that run both OSPF and BGP because OSPF converges more quickly than BGP and may begin attracting traffic before BGP has converged, resulting in dropped traffic.

If this command is configured, the router advertises its locally generated router LSAs with a metric of 0XFFFF. This action allows the router to converge but not attract transit traffic if there are better, alternative paths around this router. After the specified *announce-time* value or notification from BGP has expired, the router advertises the local router LSAs with the normal metric (interface cost).

If this command is configured with the **on-startup** keyword, then the maximum metric is temporarily set only after reboot is initiated. If this command is configured without the **on-startup** keyword, then the maximum metric is permanently used until the configuration is removed.

If the **include-stub** keyword is enabled, the stub-links in the router LSA will be sent with the max-metric. If the **summary-lsa** keyword is enabled, all self-generated summary LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter. If the **external-lsa** keyword is enabled, all self-generated external LSAs will have a metric set to 0xFF0000, unless the metric value is specified with the max-metric value parameter.

This command might be useful when you want to connect a router to an OSPF network, but do not want real traffic flowing through it if there are better, alternative paths. If there are no alternative paths, this router still accepts transit traffic as before.

Some cases where this command might be useful are as follows:

- During a router reload, you prefer that OSPF wait for BGP to converge before accepting transit traffic. If there are no alternative paths, the router still accepts transit traffic.
- A router is in critical condition (for example, it has a very high CPU load or does not have enough memory to store all LSAs or build the routing table).
- When you want to gracefully introduce or remove a router to or from the network.
- When you have a test router in a lab, connected to a production network.

**Note**

For older OSPF implementations (RFC 1247), router links in received router LSAs with a metric and cost of LSInfinity are not used during SPF calculations. Hence, no transit traffic is set to the routers originating such router LSAs.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure OSPF to originate router LSAs with the maximum metric until BGP indicates that it has converged:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# max-metric router-lsa on-startup wait-for-bgp
```

maximum interfaces (OSPF)

To limit the number of interfaces that can be configured for an Open Shortest Path First (OSPF) process, use the **maximum interfaces** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum interfaces *number-interfaces*

no maximum interfaces

Syntax Description

| | |
|--------------------------|---|
| <i>number-interfaces</i> | Number of interfaces. Range is 1 to 4294967295. |
|--------------------------|---|

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.1 | The range for number of interfaces was changed to 1 to 4294967295 from 1 to 1024. The default number of interfaces was changed to 1024 from 255. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum interface** command to increase or decrease the limit on the number of interfaces configured for an OSPF process.

You cannot configure a limit lower than the number of interfaces currently configured for the OSPF process. To lower the limit, remove interfaces from the OSPF configuration until the number of configured interfaces is at or below the desired limit. You may then apply the new, lower limit.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure a maximum interface limit of 1500 on a router:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 109  
RP/0/RSP0/CPU0:router(config-ospf)# maximum interfaces 1500
```

Related Commands

| Command | Description |
|---|--------------------------------------|
| show ospf interface, on page 1106 | Displays OSPF interface information. |

maximum paths (OSPF)

To control the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*

no maximum paths

Syntax Description

| | |
|------------------------------|---|
| <i>maximum-routes-number</i> | Maximum number of parallel routes that OSPF can install in a routing table. Range is 1 to 32. |
| Note | The maximum number of paths that can be configured is 32. |

Command Default

32 paths

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 109  
RP/0/RSP0/CPU0:router(config-ospf)# maximum paths 2
```

maximum redistributed-prefixes (OSPF)

To limit the aggregate number of prefixes that can be redistributed into an Open Shortest Path First (OSPF) process, use the **maximum redistributed-prefix** command in the appropriate mode. To return to the default limit, use the **no** form of this command.

maximum redistributed-prefixes *maximum* [*threshold-value*] [**warning-only**]

no maximum redistributed-prefixes

Syntax Description

| | |
|------------------------|---|
| <i>maximum</i> | Number of routes. Range is 1 to 4294967295. |
| <i>threshold-value</i> | (Optional) Threshold value (as a percentage) at which to generate a warning message. Range is 1 to 100. |
| warning-only | (Optional) Gives only a warning when the limit is exceeded. |

Command Default

If the command is not specified, the default is 10000.

The threshold value defaults to 75 percent.

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum redistributed-prefixes** command to increase or decrease the maximum number of prefixes (also referred to as routes) redistributed for an OSPF process.

If the *maximum* value is less than the existing number of routes, existing routes remain configured, but no new routes are redistributed.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure a maximum number of routes that can be redistributed for an OSPF routing process:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# maximum redistributed-prefixes 15000
```

Related Commands

| Command | Description |
|--|-----------------------------------|
| show ospf routes, on page 1131 | Displays the OSPF topology table. |

message-digest-key

To specify a key used with Open Shortest Path First (OSPF) Message Digest 5 (MD5) authentication, use the **message-digest-key** command in the appropriate mode. To remove an old MD5 key, use the **no** form of this command.

```
message-digest-key key-id md5 {key| clear key| encrypted key}
```

```
no message-digest-key key-id
```

Syntax Description

| | |
|------------------|--|
| <i>key-id</i> | Key number. Range is 1 to 255. |
| md5 | Enables OSPF MD5 authentication. |
| <i>key</i> | Alphanumeric string of up to 16 characters. |
| clear | Specifies that the key be clear text. |
| encrypted | Specifies that the key be encrypted using a two-way algorithm. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the message digest key parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the message digest key parameter specified for the process.

If this command is not specified at any level, then OSPF MD5 authentication is disabled.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Usually, one key individual interface is used to generate authentication information when packets are sent and to authenticate incoming packets. The same key identifier on the neighbor router must have the same *key* value.

For authentication to be enabled, you must configure the **message-digest-key** command together with the **authentication** command and its **message-digest** keyword. Both the **message-digest-key** and **authentication** commands can be inherited from a higher configuration level.

The process of changing keys is as follows. Suppose the current configuration is:

```
interface GigabitEthernet 0/3/0/2
 message-digest-key 100 md5 OLD
```

You change the configuration to the following:

```
interface GigabitEthernet 0/3/0/2
 message-digest-key 101 md5 NEW
```

The system assumes its neighbors do not have the new key yet, so it begins a rollover process. It sends multiple copies of the same packet, each authenticated by different keys. In this example, the system sends out two copies of the same packet—the first one authenticated by key 100 and the second one authenticated by key 101.

Rollover allows neighboring routers to continue communication while the network administrator is updating them with the new key. Rollover stops after the local system finds that all its neighbors know the new key. The system detects that a neighbor has the new key when it receives packets from the neighbor authenticated by the new key.

After all neighbors have been updated with the new key, the old key should be removed. In this example, you would enter the following:

```
interface ethernet 1
 no ospf message-digest-key 100
```

Then, only key 101 is used for authentication on interface 1.

We recommend that you not keep more than one key individual interface. Every time you add a new key, you should remove the old key to prevent the local system from continuing to communicate with a hostile system that knows the old key. Removing the old key also reduces overhead during rollover.

**Note**

The MD5 key is always stored in encrypted format on the router. The **clear** and **encrypted** keywords inform the router whether the value that is entered is encrypted or unencrypted.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set a new key 19 with the password *8ry4222* :

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0

RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# message-digest-key 19 md5 8ry4222
```

Related Commands

| Command | Description |
|--|---|
| area (OSPF), on page 934 | Configures an OSPF area. |
| authentication (OSPF), on page 936 | Enables plain text, MD5 authentication, or null authentication for an OSPF interface. |
| default-cost (OSPF), on page 962 | Enables authentication for an OSPF area. |

mpls ldp auto-config (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) interface automatic configuration, use the **mpls ldp auto-config** command in the appropriate mode. To disable LDP-IGP interface automatic configuration, use the **no** form of this command.

mpls ldp auto-config

no mpls ldp auto-config

Command Default LDP-IGP interface automatic configuration is disabled for OSPF.

Command Modes Interface configuration
Area configuration
Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to enable LDP-IGP interface automatic configuration:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 01
RP/0/RSP0/CPU0:router(config-ospf)# mpls ldp auto-config
```

mpls ldp sync (OSPF)

To enable Label Distribution Protocol (LDP)-Interior Gateway Protocol (IGP) synchronization, use the **mpls ldp sync** command in the appropriate mode. To disable LDP-IGP synchronization, use the **no** form of this command.

mpls ldp sync [disable]

no mpls ldp sync

Syntax Description

| | |
|----------------|--|
| disable | (Optional) Disables MPLS LDP synchronization from within the OSPF interface and area configuration submodes only. For the OSPF router configuration mode, use the no form of the command. |
|----------------|--|

Command Default

LDP-IGP synchronization is disabled for OSPF.

Command Modes

Interface configuration

Area configuration

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable LDP-IGP synchronization:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 01
RP/0/RSP0/CPU0:router(config-ospf)# mpls ldp sync
```

mpls traffic-eng (OSPF)

To configure an Open Shortest Path First (OSPF) area for Multiprotocol Label Switching traffic engineering (MPLS TE), use the **mpls traffic-eng** command in the appropriate configuration mode. To remove the MPLS TE from an area, use the **no** form of this command.

mpls traffic-eng

no mpls traffic-eng

Syntax Description This command has no keywords or arguments.

Command Default MPLS TE is not configured for OSPF.

Command Modes Area configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **mpls traffic-eng** command for OSPF to support MPLS traffic engineering. OSPF provides the flooding mechanism that is used to flood TE link information.



Note This command is supported only in the default VRF mode.

We recommend that you configure the **mpls traffic-eng router-id** command instead of using the **router-id** command in global configuration mode.

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to associate loopback interface 0 with area 0, and area 0 is declared to be an MPLS area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# router-id 10.10.10.10
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface loopback 0
```

Related Commands

| Command | Description |
|--|--|
| capability opaque disable, on page 943 | Controls the OSPF opaque LSA support capability. |
| mpls traffic-eng multicast-intact (OSPF), on page 1022 | Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface. |
| router-id (OSPF), on page 1076 | Configures a router ID for the OSPF process. |

mpls traffic-eng igp-intact (OSPF)

To ensure that the OSPF protocol installs at least one IPv4 next-hop when it adds the tunnel next-hops (igp-shortcuts), use the **mpls traffic-eng igp-intact** command in the router configuration mode. To disable IGP-intact, use the **no** form of this command.

mpls traffic-eng igp-intact

no mpls traffic-eng igp-intact

Command Default IGP-intact is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The OSPF protocol adds both tunnel next-hops (igp-shortcuts) and IPv4 next-hops to the next-hop list in the Routing Information Base (RIB), until it reaches the maximum number of paths. When IGP-intact is enabled, it ensures that the Routing Information Base (RIB) always has at least one IPv4 next-hop present in the list of next-hops when the number of paths is at maximum.



Note IGP-intact should be used only when Policy-Based Tunnel Selection (PBTS) is in use.

Configure IGP-intact only when Policy-Based Tunnel Selection (PBTS) is in use. This ensures that at least one IPv4 next-hop is available for the default Differentiated Services Code Point (DSCP) traffic class. It also prevents traffic loss for other DSCP traffic classes by diverting such traffic to an IPv4 next-hop when the corresponding tunnel is unavailable for forwarding.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to enable IGP-intact:

```
RP/0/RSP0/CPU0:router# configure
```

```
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng igp-intact
```

Related Commands

| Command | Description |
|--|---|
| maximum paths (OSPF), on page 1009 | Configures the maximum number of parallel routes that the Open Shortest Path First (OSPF) protocol can support. |

mpls traffic-eng multicast-intact (OSPF)

To enable multicast-intact for Open Shortest Path First (OSPF) routes so that multicast-intact paths are published to the Routing Information Base (RIB), use the **mpls traffic-eng multicast-intact** command in the appropriate configuration mode. To remove the MPLS TE area, use the **no** form of this command.

mpls traffic-eng multicast-intact

no mpls traffic-eng multicast-intact

Syntax Description This command has no keywords or arguments.

Command Default MPLS TE is not configured for OSPF.

Command Modes Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF support for MPLS TE is a component of the overall MPLS TE feature. Other MPLS TE software components must also be configured for this feature to be fully supported.



Note

This command is supported only in the default VRF mode.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable publishing of multicast-intact paths to RIB:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng multicast-intact
```

mpls traffic-eng router-id (OSPF)

To specify that the traffic engineering router identifier for the node is the IP address associated with a given Open Shortest Path First (OSPF) interface, use the **mpls traffic-eng router-id** command in the appropriate configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id {*router-id*} *type interface-path-id*

no mpls traffic-eng router-id {*router-id*} *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>router-id</i> | The 32-bit router ID value specified in four-part, dotted-decimal notation (must be in the valid IP address range of 0.0.0.0 to 255.255.255.255). |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

If this command is specified in router configuration mode, then the traffic engineering router identifier for the node is the IP address associated with a given interface.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

This identifier of the router acts as a stable IP address for the traffic engineering configuration. This IP address is flooded to all nodes. For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router identifier of the destination node, because that is the address that the traffic engineering topology database at the tunnel head uses for its path calculation.

**Note**

We recommend that loopback interfaces be used for Multiprotocol Label Switching traffic engineering (MPLS TE), because they are more stable than physical interfaces.

**Note**

This command is supported only in the default VRF mode.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify the traffic engineering router identifier as the IP address associated with loopback interface 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# mpls traffic-eng router-id loopback 0
```

Related Commands

| Command | Description |
|---|--------------------------------------|
| mpls traffic-eng (OSPF), on page 1018 | Configures an OSPF area for MPLS TE. |

mtu-ignore (OSPF)

To prevent Open Shortest Path First (OSPF) from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in the appropriate mode. To reset to default, use the **no** form of this command.

mtu-ignore [**disable**| **enable**]

no mtu-ignore

Syntax Description

| | |
|----------------|--|
| disable | (Optional) Enables checking for whether OSPF neighbors are using the MTU on a common interface. |
| enable | (Optional) Disables checking for whether OSPF neighbors are using the MTU on a common interface. |

Command Default

The default is **mtu-ignore** with no keywords, which disables MTU checking.

If this command is not specified in interface configuration mode, then the interface adopts the MTU ignore parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the MTU ignore parameter specified for the process.

If this command is not specified at any level, then OSPF checks the MTU received from neighbors when exchanging DBD packets.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration
 Multi-area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF checks whether OSPF neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency is not established.

The keywords, **disable** and **enable**, do not need to be used. If no keywords are used, the **mtu-ignore** command disables MTU checking. You can then use the **no mtu-ignore** command to activate MTU checking.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to disable MTU mismatch detection on receiving DBD packets:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# mtu-ignore
```

multi-area-interface

To enable multiple adjacencies for different Open Shortest Path First (OSPF) areas and enter multi-area interface configuration mode, use the **multi-area-interface** command in the area configuration mode. To reset to the default, use the **no** form of this command.

multi-area-interface *type interface-path-id*

no multi-area-interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

An OSPF network is enabled for one area only.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **multi-area-interface** command to enable area border routers (ABRs) to establish multiple adjacencies for different OSPF areas.

Each multiple area adjacency is announced as a point-to-point unnumbered link in the configured area. This point-to-point link provides a topological path for that area. The first or primary adjacency using the link advertises the link consistent with draft-ietf-ospf-multi-area-adj-06.txt.

You can configure multi-area adjacency on any interface where only two OSPF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OSPF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable multiple area adjacency for OSPF 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# multi-area-interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-mif)# ?

authentication          Enable authentication
authentication-key      Authentication password (key)
commit                  Commit the configuration changes to running
cost                    Interface cost
database-filter          Filter OSPF LSA during synchronization and flooding
dead-interval            Interval after which a neighbor is declared dead
describe                Describe a command without taking real actions
distribute-list          Filter networks in routing updates
do                       Run an exec command
exit                     Exit from this submode
hello-interval           Time between HELLO packets
message-digest-key       Message digest authentication password (key)
mtu-ignore               Enable/Disable ignoring of MTU in DBD packets
no                       Negate a command or set its defaults
packet-size              Customize size of OSPF packets upto MTU
pwd                      Commands used to reach current submode
retransmit-interval      Time between retransmitting lost link state advertisements
root                     Exit to the global configuration mode
show                     Show contents of configuration
transmit-delay           Estimated time needed to send link-state update packet
RP/0/RSP0/CPU0:router(config-ospf-ar-mif)#
```

Related Commands

| Command | Description |
|---|--------------------------------------|
| show ospf interface, on page 1106 | Displays OSPF interface information. |

neighbor (OSPF)

To configure Open Shortest Path First (OSPF) routers interconnecting to nonbroadcast networks, use the **neighbor** command in interface configuration mode. To remove a configuration, use the **no** form of this command.

neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]

no neighbor *ip-address* [**cost** *number*] [**priority** *number*] [**poll-interval** *seconds*]

Syntax Description

| | |
|-------------------------------------|--|
| <i>ip-address</i> | Interface IP address of the neighbor. |
| cost <i>number</i> | (Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the cost command. On point-to-multipoint interfaces, cost <i>number</i> is the only keyword and argument combination that works. The cost keyword does not apply to nonbroadcast multiaccess (NBMA) networks. |
| priority <i>number</i> | (Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The priority keyword does not apply to point-to-multipoint interfaces. |
| poll-interval <i>seconds</i> | (Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The poll-interval keyword does not apply to point-to-multipoint interfaces. |

Command Default

No configuration is specified.

priority *number* : 0

poll-interval *seconds* : 120 seconds (2 minutes)

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must include one neighbor entry in the software configuration for each known nonbroadcast network neighbor. The neighbor address must be on the primary address of the interface.

If a neighboring router has become inactive (hello packets have not been received for the router dead interval period), it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at a reduced rate called the *poll interval*.

When the router starts up, it sends only hello packets to those routers with nonzero priority; that is, routers that are eligible to become designated routers (DRs) and backup designated routers (BDRs). After the DR and BDR are selected, the DR and BDR start sending hello packets to all neighbors to form adjacencies.

To filter all outgoing OSPF link-state advertisement (LSA) packets for the neighbor, use the **neighbor database-filter all out** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to declare a router at address 172.16.3.4 on a nonbroadcast network, with a priority of 1 and a poll interval of 180 seconds:

```
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
```

The following example illustrates a network with nonbroadcast:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router(config-if)# ip address 172.16.3.10 255.255.255.0

RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network nonbroadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval 180
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.5 cost 10 priority 1 poll-interval 180
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.6 cost 15 priority 1 poll-interval 180
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.7 priority 1 poll-interval 180
```

Related Commands

| Command | Description |
|---|---|
| neighbor database-filter all out , on page 1033 | Filters all outgoing LSAs to an OSPF neighbor. |
| network (OSPF) , on page 1035 | Configures the OSPF network type to a type other than the default for a given medium. |
| priority (OSPF) , on page 1052 | Sets the router priority, which helps determine the designated router for this network. |

neighbor database-filter all out

To filter all outgoing link-state advertisements (LSAs) to an Open Shortest Path First (OSPF) neighbor, use the **neighbor database-filter all out** command in interface configuration mode. To restore the forwarding of LSAs to the neighbor, use the **no** form of this command.

neighbor ip-address database-filter all out

no neighbor ip-address database-filter all out

Syntax Description

| | |
|-------------------|--|
| <i>ip-address</i> | IP address of the neighbor to which outgoing LSAs are blocked. |
|-------------------|--|

Command Default

Instead of all outgoing LSAs being filtered to the neighbor, they are flooded to the neighbor.

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **neighbor database-filter all out** command to filter all outgoing OSPF LSA packets during synchronization and flooding for point-to-multipoint neighbors on nonbroadcast networks. More neighbor options are available with the **neighbor** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to prevent flooding of OSPF LSAs from point-to-multipoint networks to the neighbor at IP address 10.2.3.4:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.2.3.4 database-filter all out
```

Related Commands

| Command | Description |
|--|---|
| neighbor (OSPF) , on page 1030 | Configures OSPF routers interconnecting to nonbroadcast networks. |

network (OSPF)

To configure the Open Shortest Path First (OSPF) network type to a type other than the default for a given medium, use the **network** command in the appropriate mode. To return to the default value, use the **no** form of this command.

```
network {broadcast| non-broadcast| {point-to-multipoint [non-broadcast]| point-to-point}}
no network
```

Syntax Description

| | |
|----------------------------|---|
| broadcast | Sets the network type to broadcast. |
| non-broadcast | Sets the network type to nonbroadcast multiaccess (NBMA). |
| point-to-multipoint | Sets the network type to point-to-multipoint. |
| non-broadcast | (Optional) Sets the point-to-multipoint network to be nonbroadcast. If you use this keyword, the neighbor command is required. |
| point-to-point | Sets the network type to point-to-point. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the network parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the network parameter specified for the process.

If this command is not specified at any level, then the OSPF network type is the default of the given medium. GigabitEthernet and TenGigEthernet interfaces default to broadcast.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Configuring NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPF network type as a point-to-multipoint network. Routing between two routers that are not directly connected go through the router that has virtual circuits to both routers.

If this command is issued on an interface that does not allow it, this command is ignored.

OSPF has two features related to point-to-multipoint networks. One feature applies to broadcast networks; the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the OSPF network as a nonbroadcast network:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 172.16.3.4 priority 1 poll-interval
180
```

Related Commands

| Command | Description |
|--|---|
| neighbor (OSPF) , on page 1030 | Configures OSPF routers interconnecting to nonbroadcast networks. |

nsf (OSPF)

To configure nonstop forwarding (NSF) for the Open Shortest Path First (OSPF) protocol, use the **nsf** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no nsf** form of this command.

```
nsf {cisco [enforce global]| ietf [helper disable]}
```

```
no nsf {cisco [enforce global]| ietf [helper disable]}
```

Syntax Description

| | |
|-----------------------|--|
| cisco | Enables Cisco Nonstop Forwarding. |
| enforce global | (Optional) Cancels NSF restart when non-NSF network device neighbors are detected. |
| ietf | Enables Internet Engineering Task Force (IETF) graceful restart. |
| helper disable | (Optional) Disables router helper support. |

Command Default

NSF is disabled.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The NSF feature allows for the forwarding of data packets to continue along known routes while routing protocol information (such as OSPF) is being restored following a switchover.

Use the **nsf** command if the router is expected to perform NSF during restart. To experience the full benefits of this feature, configure all neighboring routers with NSF.

When this command is used without the optional **cisco enforce global** keywords and non-NSF neighbors are detected, the NSF restart mechanism aborts on the interfaces of those neighbors and functions properly on others.

When this command is used with the optional **cisco enforce global** keywords and non-NSF neighbors are detected, NSF restart is canceled for the entire OSPF process.

IETF graceful restart provides an NSF mechanism to allow data traffic to flow seamlessly with no packet drops during the transient period when OSPF attempts to recover after a process restart or RP failover, within the guidelines of RFC 3623.

By default, neighbors in helper mode listen to both the NSF Cisco- and NSF IETF-type LSAs. The **nsf** command enables one type of mechanism that would undergo an RP failover or, anticipating an OSPF process restart. If the **cisco** or **ietf** keyword is not entered, NSF is not enabled, irrespective of neighbors in listening mode for both NSF Cisco and NSF IETF.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to cancel NSF restart for the entire OSPF process if non-NSF neighbors are detected on any network interface during restart:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf cisco enforce global
```

nsf flush-delay-time (OSPF)

To configure the maximum time allowed for nonstop forwarding (NSF) external route queries for the Open Shortest Path First (OSPF) protocol, use the **nsf flush-delay-time** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf flush-delay-time *seconds*

no nsf flush-delay-time *seconds*

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Length of time (in seconds) allowed for NSF external route queries. Range is 1 to 3600 seconds. |
|----------------|---|

Command Default

seconds : 300

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the maximum time for NSF to learn external routes for OSPF at 60 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsf flush-delay-time 60
```

nsf interval (OSPF)

To configure the minimum time between consecutive nonstop forwarding (NSF) restart attempts for the Open Shortest Path First (OSPF) protocol, use the **nsf interval** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf interval *seconds*

no nsf interval *seconds*

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Length of time (in seconds) between consecutive restart attempts. Range is 90 to 3600 seconds. |
|----------------|--|

Command Default

seconds : 90

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use the **nsf interval** command, the OSPF process must be up for at least 90 seconds before OSPF attempts to perform an NSF restart.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the minimum time between consecutive NSF restart attempts at 120 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf)# nsf interval 120
```

nsf lifetime (OSPF)

To configure the maximum time that routes are held in the Routing Information Base (RIB) following an Open Shortest Path First (OSPF) process restart, use the **nsf lifetime** command in the appropriate mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf lifetime *seconds*

no nsf lifetime *seconds*

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | The length of time (in seconds) that routes are held in the RIB. Range is 90 to 3600 seconds. |
|----------------|---|

Command Default

seconds : 95

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use this command, the OSPF process must reconverge within the maximum length of time configured. If the convergence exceeds this length of time, routes are purged from RIB and nonstop forwarding (NSF) restart may fail.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the maximum lifetime for OSPF NSF at 120 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
```

```
RP/0/RSP0/CPU0:router(config-ospf) # nsf lifetime 120
```

nsr (OSPF)

To configure nonstop routing (NSR) for the Open Shortest Path First (OSPF) protocol, use the **nsr** command in OSPF router configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no nsr** form of this command.

nsr

no nsr

Command Default NSR is not defined.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The NSR feature allows an OSPF process on the active RP to synchronize all necessary data and states with the OSPF process on the standby RP. When the switchover happens, the OSPF process on the newly active RP has all the necessary data and states to continue running and does not require any help from its neighbors.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to configure NSR:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# nsr
```

nssa (OSPF)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*] [**metric-type** *type-value*]] [**no-summary**]

no nssa

Syntax Description

| | |
|--------------------------------------|---|
| no-redistribution | (Optional) Imports routes only into the normal areas, but not into the NSSA area, by the redistribute command when the router is an NSSA Area Border Router (ABR). |
| default-information-originate | (Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA Autonomous System Boundary Router (ASBR). |
| metric <i>metric-value</i> | (Optional) Specifies the metric used for generating the default route. If you omit a value and do not specify a value using the defaultmetric command, the default metric value is 10. Range is 1 to 16777214. |
| metric-type <i>type-value</i> | (Optional) Specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of the following values: 1—Type 1 external route 2—Type 2 external route |
| no-summary | (Optional) Prevents an ABR from sending summary link advertisements into the NSSA. |

Command Default

No NSSA area is defined.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An NSSA does not flood Type 5 external LSAs from the core into the area, but can import autonomous system external routes in a limited fashion within the area.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows how to configure area 1 as an NSSA area:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 1
RP/0/RSP0/CPU0:router(config-ospf-ar)# nssa
```

ospf name-lookup

To configure the Open Shortest Path First (OSPF) protocol to look up Domain Name System (DNS) names, use the **ospf name-lookup** command in global configuration mode. To disable this function, use the **no** form of this command.

ospf name-lookup

no ospf name-lookup

Command Default Routers are displayed by router ID or neighbor ID.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ospf name-lookup** command to easily identify a router when executing all OSPF **show** command displays. The router is displayed by name rather than by its router ID or neighbor ID.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to configure OSPF to identify a router by name:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# ospf name-lookup
```

packet-size (OSPF)

To configure the size of Open Shortest Path First (OSPF) packets up to the size specified by the maximum transmission unit (MTU), use the **packet-size** command in the appropriate configuration mode. To disable this function and reestablish the default packet size, use the **no** form of this command.

packet-size *bytes*

no packet-size

Syntax Description

| | |
|--------------|--|
| <i>bytes</i> | Size, in bytes. Range is 576 to 10000 bytes. |
|--------------|--|

Command Default

If the command is not specified, the default packet size is either the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes.

Command Modes

Router configuration
 Area configuration
 Interface configuration
 VRF configuration
 Multi-area configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | The default packet size was changed to the lower interface IP MTU size or 9000 bytes. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **packet-size** command to customize the size of OSPF packets. The OSPF protocol compares the packet size and the MTU size and uses the lower packet size value.

If the command is not configured, the default packet size is equal to the interface IP MTU size (if that is lower than 9000 bytes) or 9000 bytes. For example, if the interface IP MTU size is 1500 bytes, OSPF uses packet size of 1500 bytes on the interface because the byte size is lower than 9000 bytes. If the interface IP MTU size is 9500 bytes, OSPF uses packet size of 9000 bytes on the interface because the byte size exceeds 9000 bytes. The interface IP MTU size depends on the interface and the platform. In most cases, the default interface IP MTU value will be lower than 9000 bytes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to configure the packet size on an interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 1/0/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# packet-size 3500
```

passive (OSPF)

To suppress the sending of Open Shortest Path First (OSPF) protocol operation on an interface, use the **passive** command in the appropriate mode. To remove the passive configuration, use the **no** form of this command.

passive [**disable**| **enable**]

no passive

Syntax Description

| | |
|----------------|---|
| disable | (Optional) Sends OSPF updates. |
| enable | (Optional) Disables sending OSPF updates. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPF updates are sent on the interface.

Command Modes

Interface configuration
Area configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF routing information is neither sent nor received through the specified interface. The interface appears as a stub network in the OSPF router (Type 1) link-state advertisement (LSA).

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows that GigabitEthernet interface 1/0/0/2 reduces OSPF updates because passive mode is enabled; however, GigabitEthernet interface 0/1/0/3 receives normal OSPF traffic flow:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router (config)# router ospf 1
RP/0/RSP0/CPU0:router (config-ospf)# area 0
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet1/0/0/2
RP/0/RSP0/CPU0:router (config-ospf-ar-if)# passive
RP/0/RSP0/CPU0:router (config-ospf-ar-if)# exit
RP/0/RSP0/CPU0:router (config-ospf-ar)# interface GigabitEthernet1/0/0/3
RP/0/RSP0/CPU0:router (config-ospf-ar-if)# end
```

priority (OSPF)

To set the router priority for an interface, which helps determine the designated router for an Open Shortest Path First (OSPF) link, use the **priority** command in the appropriate mode. To return to the default value, use the **no** form of this command.

priority *value*

no priority *value*

Syntax Description

| | |
|--------------|---|
| <i>value</i> | 8-bit unsigned integer indicating the router priority value. Range is 0 to 255. |
|--------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified for the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration
Area configuration
Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to multiaccess networks (in other words, not point-to-point networks).

This priority value is used when you configure the Open Shortest Path First (OSPF) protocol for nonbroadcast networks using the **neighbor** command for OSPF.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows that priority is set through the **priority** and **neighbor** commands for Routers A and B and that the neighbor priority value must reflect that of the neighbor router:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.0.0.2 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# priority 4
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.1 priority 6

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet POS 0/2/0/1
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.0.0.1 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# priority 6
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# neighbor 10.0.0.2 priority 4
```

Related Commands

| Command | Description |
|---|---|
| neighbor (OSPF), on page 1030 | Configures OSPF routers interconnecting to nonbroadcast networks. |
| network (OSPF), on page 1035 | Configures the OSPF network type to a type other than the default for a given medium. |

protocol shutdown

To disable an instance of the Open Shortest Path First (OSPF) protocol so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the router configuration mode. To reenble the OSPF protocol, use the **no** form of this command.

protocol shutdown

no protocol shutdown

Command Default No default behavior or values

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **protocol shutdown** command to disable the OSPF protocol for a specific routing instance without removing any existing OSPF configuration parameters.

The OSPF protocol continues to run on the router and you can use the current OSPF configuration, but OSPF does not form any adjacencies on any interface.

This command is similar to performing the **no router ospf** command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to disable the OSPF 1 instance:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# protocol shutdown
```

queue dispatch flush-lsa

To change the number of LSAs scheduled (rate-limited) for flushing, that are processed in each iteration, use the **queue dispatch flush-lsa** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch flush-lsa *count*

no queue dispatch flush-lsa

Syntax Description

| | |
|--------------|--|
| <i>count</i> | Maximum number of LSAs flushed per run. Range is 30 to 3000. |
|--------------|--|

Command Default

The default LSAs flushed per run is 150 (when the count is not configured).

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to limit the number of LSAs flushed per run to 30:

```
RP/0/RSP0/CPU0:router (config-ospf) # queue dispatch flush-lsa 30
```

Use the [show ospf message-queue](#), [on page 1114](#) command to see the queue dispatch values, peak lengths, and limits.

Related Commands

| Command | Description |
|---|---|
| queue dispatch incoming, on page 1057 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 1059 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 1061 | Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run. |
| queue limit, on page 1063 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 1114 | Displays the information about the queue dispatch values, peak lengths, and limits. |

queue dispatch incoming

To limit the number of incoming packets (LSAUpdates, LSAsacks, DBDs, LSRequests, and Hellos that trigger a change state) processed, use the **queue dispatch incoming** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch incoming *count*

no queue dispatch incoming

Syntax Description

| | |
|--------------|---|
| <i>count</i> | Maximum number of continuous events processed. Range is 30 to 3000. |
|--------------|---|

Command Default

The default incoming count is 300 packets (when the count is not configured).

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how limit the number of incoming packets processed to 500:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch incoming 500
```

Use the [show ospf message-queue](#), on page 1114 command to see the queue dispatch values, peak lengths, and limits.

Related Commands

| Command | Description |
|---|---|
| queue dispatch rate-limited-lsa, on page 1059 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 1061 | Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run. |
| queue limit, on page 1063 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 1114 | Displays the information about the queue dispatch values, peak lengths, and limits. |

queue dispatch rate-limited-lsa

To set the maximum number of rate-limited link-state advertisement (LSA) (re-)originations processed per run, use the **queue dispatch rate-limited-lsa** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch rate-limited-lsa *count*

no queue dispatch rate-limited-lsa

Syntax Description

| | |
|--------------|---|
| <i>count</i> | Maximum number of rate-limited LSAs processed per run. Range is 30 to 3000. |
|--------------|---|

Command Default

The default number of rate-limited LSAs processed per run is 300 (when this count is not configured).

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the maximum number of rate-limited LSA (re-)originations processed per run to 300:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch rate-limited-lsa 300
```

Related Commands

| Command | Description |
|--|---|
| queue dispatch incoming, on page 1057 | Limits the number of continuous incoming events processed. |
| queue dispatch spf-lsa-limit, on page 1061 | Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run. |
| queue limit, on page 1063 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 1114 | Displays the information about the queue dispatch values, peak lengths, and limits. |

queue dispatch spf-lsa-limit

To change the maximum number of Type 3-4 and Type 5-7 link-state advertisements (LSAs) processed per shortest path first (SPF) iteration within a single SPF run, use the **queue dispatch spf-lsa-limit** command in router configuration mode. To return to the system default value, use the **no** form of this command.

queue dispatch spf-lsa-limit *count*

no queue dispatch spf-lsa-limit

Syntax Description

| | |
|--------------|---|
| <i>count</i> | Maximum number of continuous Type 3-4 and Type 5-7 LSAs processed per SPF in each scheduled iteration within a single SPF run. Range is 30 to 3000. |
|--------------|---|

Command Default

The default number of Type 3-4 and Type 5-7 processed per run is 150 LSAs (when this command is not configured).

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to limit the number of continuous Type 3-4 and Type 5-7 LSAs processed by SPF per scheduling run, to 100:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue dispatch spf-lsa-limit 100
```

Related Commands

| Command | Description |
|---|--|
| queue dispatch incoming, on page 1057 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 1059 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run |
| queue limit, on page 1063 | Sets the high watermark for incoming priority events. |
| show ospf message-queue, on page 1114 | Displays the information about the queue dispatch values, peak lengths, and limits. |

queue limit

To set the high watermark for incoming events by priority, use the **queue limit** in router configuration mode. To return to the system default values, use the **no** form of this command.

queue limit {**high**| **medium**| **low**} *count*

no queue limit {**high**| **medium**| **low**}

Syntax Description

| | |
|---------------|---|
| high | High watermark for incoming high-priority events (state-changing Hellos). |
| medium | High watermark for incoming medium-priority events (LSA ACK). |
| low | High watermark for incoming low-priority events (DBD/LSUpd/LSReq). |
| <i>count</i> | Maximum number of events per queue. Events are dropped when the priority queue size exceeds this value. Range is 1000 to 30000. |

Command Default

High watermark: 9500 (when the corresponding configuration is not present).

Medium watermark: 9000 (when the corresponding configuration is not present).

Low watermark: 8000 (when the corresponding configuration is not present).

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Always keep the limits in the following order of priority:

Limit for High > Limit for Medium > Limit for Low

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following examples show how to set the maximum number of events per queue:

```
RP/0/RSP0/CPU0:router(config-ospf)# queue limit high 11000
RP/0/RSP0/CPU0:router(config-ospf)# queue limit medium 10000
RP/0/RSP0/CPU0:router(config-ospf)# queue limit low 9000
```

Related Commands

| Command | Description |
|---|---|
| queue dispatch incoming, on page 1057 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 1059 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 1061 | Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run. |
| show ospf message-queue, on page 1114 | Displays the information about the queue dispatch values, peak lengths, and limits. |

range (OSPF)

To consolidate and summarize routes at an area boundary, use the **range** command in area configuration mode. To disable this function, use the **no** form of this command.

range *ip-address mask* [**advertise**|**not-advertise**]

no range *ip-address mask* [**advertise**|**not-advertise**]

Syntax Description

| | |
|----------------------|--|
| <i>ip-address</i> | IP address in four-part, dotted-decimal notation. |
| <i>mask</i> | IP address mask. |
| advertise | (Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA). |
| not-advertise | (Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks. |

Command Default

When this command is not specified for Area Border Routers (ABRs), routes at an area boundary are not consolidated or summarized.

Advertise is the default.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **range** command only with Area Border Router (ABRs). Use the command to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This process is called *route summarization*.

Multiple **range** configurations specifying the **range** command can be configured. Thus, the OSPF protocol can summarize addresses for many different sets of address ranges.

The summarized route uses the maximum cost of the routes assumed in the range.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows area 36.0.0.0 consisting of interfaces whose IP addresses have “10.31.x.x” as the first two octets. The **range** command summarizes interfaces. Instead of advertising eight networks individually, the single route 10.31.0.0 255.255.0.0 is advertised:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet 0/3/0/2
!
RP/0/RSP0/CPU0:router(config-ospf)# area 36.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# range 10.31.0.0 255.255.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/1/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/2
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# interface GigabitEthernet0/2/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# end
```

Related Commands

| Command | Description |
|---|--|
| summary-prefix (OSPF), on page 1155 | Creates aggregate addresses for routes being redistributed from another routing protocol into the OSPF protocol. |

redistribute (OSPF)

To redistribute routes from one routing domain into Open Shortest Path First (OSPF), use the **redistribute** command in the appropriate mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)

redistribute bgp *process-id* [**preserve-med**] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute bgp *process-id* [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Local Interface Routes

redistribute connected [**instance** *instance-name*] [**instance IPCP**][**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute connected [**instance** *instance-name*] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Directed-attached gateway redundancy (DAGR)

redistribute dagr [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute dagr [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Enhanced Interior Gateway Routing Protocol (EIGRP)

redistribute eigrp *process-id* [**match** {**external** [1|2]|**internal**}] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute eigrp *process-id* [**match** {**external** [1|2]|**internal**}] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Intermediate System-to-Intermediate System (IS-IS)

redistribute isis *process-id* [**level-1**|**level-2**|**level-1-2**] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute isis *process-id* [**level-1**|**level-2**|**level-1-2**] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Open Shortest Path First (OSPF)

redistribute ospf *process-id* [**match** {**external** [1|2]|**internal**|**nssa-external** [1|2]}] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

no redistribute ospf *process-id* [**match** {**external** [1|2]|**internal**|**nssa-external** [1|2]}] [**metric** *metric-value*] [**metric-type** {1|2}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Routing Information Protocol (RIP)

redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute rip [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

IP Static Routes

redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]
no redistribute static [**metric** *metric-value*] [**metric-type** {**1** | **2**}] [**route-policy** *policy-name*] [**tag** *tag-value*]

Syntax Description

| | |
|--|--|
| bgp | Distributes routes from the BGP protocol. |
| <i>process-id</i> | For the bgp keyword, an autonomous system number has the following ranges: <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. <p>For the isis keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> <p>For the ospf keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| preserve-med | (Optional) Preserves the Multi Exit Discriminator (MED) of BGP routes. |
| metric <i>metric-value</i> | (Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the source protocol. |
| metric-type { 1 2 } | (Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values: <ul style="list-style-type: none"> • 1 —Type 1 external route • 2 —Type 2 external route |
| tag <i>tag-value</i> | (Optional) Specifies the value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the external LSAs. Range is 0 to 4294967295. |
| route-policy <i>policy-name</i> | (Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF. |
| connected | Distributes routes that are established automatically by virtue of having enabled IP on an interface. |

| | |
|--|--|
| instance | Connected instance. |
| <i>instance-name</i> | Name of the connected instance. |
| instance IPCP | Distributes routes from IPCP protocols. |
| eigrp | Distributes routes from the EIGRP protocol. |
| isis | Distributes routes from the IS-IS protocol. |
| level-1 | (Optional) Redistributes Level 1 routes into other IP routing protocols independently. |
| level-1-2 | (Optional) Distributes both Level 1 and Level 2 routes into other IP routing protocols. |
| level-2 | (Optional) Distributes Level 2 routes into other IP routing protocols independently. |
| ospf | Distributes routes from the OSPF protocol. |
| match { internal external [1 2] nssa-external [1 2] } | <p>(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:</p> <ul style="list-style-type: none"> • internal—Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes). • external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes. • nssa-external [1 2]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes. <p>For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.</p> <p>If no match is specified, the default is no filtering.</p> |
| rip | Distributes routes from the RIP protocol. |
| static | Distributes IP static routes. |
| dagr | Distributes routes from the directed-attached gateway redundancy (DAGR). |

Command Default

Route redistribution is disabled.

metric *metric-value*: Default is 20 for routes from all protocols except BGP routes, for which the default is 1.

metric-type : Type 2 external route.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. Support was added for redistribution from directed-attached gateway redundancy (DAGR). The keyword dagr was added. The instance keyword and <i>instance-name</i> argument were added for connected routes. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

When redistributing routes (into OSPF) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Redistributed routing information should always be filtered by the **policy** *policy-name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed into OSPF.

For information about routing policies, see the *Routing Policy Commands on the Cisco ASR 9000 Series Router* module of *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Whenever you use the **redistribute** or [default-information originate \(OSPF\)](#), on page 964 command to redistribute routes into an OSPF routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPF routing domain.

When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to cause BGP routes to be redistributed into an OSPF domain:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 110
RP/0/RSP0/CPU0:router(config-ospf)# redistribute bgp 100
```

The following example shows how to redistribute the specified IS-IS process routes into an OSPF domain. The IS-IS routes are redistributed with a metric of 100.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# redistribute isis 108 metric 100
```

In the following example, network 10.0.0.0 appears as an external link-state advertisement (LSA) in OSPF 1:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-if)# ip address 10.0.0.0 255.0.0.0
!
RP/0/RSP0/CPU0:router(config)# interface GigabitEthernet 0/2/0/2
RP/0/RSP0/CPU0:router(config)# ip address 10.99.0.0 255.0.0.0
!
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# redistribute ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/2
!
RP/0/RSP0/CPU0:router(config)# router ospf 2
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/1/0/1
```

Related Commands

| Command | Description |
|--|---|
| default-information originate (OSPF) , on page 964 | Generates a default external route into an OSPF routing domain. |

retransmit-interval (OSPF)

To specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the Open Shortest Path First (OSPF) interface, use the **retransmit-interval** command in the appropriate mode. To return to the default value, use the **no** form of this command.

retransmit-interval *seconds*

no retransmit-interval

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds. |
|----------------|--|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified for the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to set the retransmit interval value to 8 seconds in interface configuration mode:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# retransmit-interval 8
```

route-policy (OSPF)

To specify a routing policy to filter Type 3 link-state advertisements (LSA), use the **route-policy** command in area configuration mode. To disable the routing policy, use the **no** form of this command.

route-policy *route-policy-name* {**in**|**out**}

no route-policy *route-policy-name* {**in**|**out**}

Syntax Description

| | |
|--------------------------|------------------------------------|
| <i>route-policy-name</i> | Name of route policy. |
| in | Applies policy to inbound routes. |
| out | Applies policy to outbound routes. |

Command Default

No policy is applied.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify an OSPF routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify an OSPF route policy for inbound routes in area 0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-area)# route-policy area0_in in
```


router-id (OSPF)

To configure a router ID for the Open Shortest Path First (OSPF) process, use the **router-id** command in the appropriate mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command after clearing or restarting the OSPF process.

router-id router-id

no router-id router-id

Syntax Description

| | |
|------------------|---|
| <i>router-id</i> | 32-bit router ID value specified in four-part, dotted-decimal notation. |
|------------------|---|

Command Default

If this command is not configured, the router ID is the highest IP version 4 (IPv4) address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

It is good practice to use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This action ensures that OSPF can function regardless of the interface address configuration. Clear the OSPF process using the **clear ospf process** command or restart the OSPF process for the **no router-id** command to take effect.

OSPF attempts to obtain a router ID in the following ways (in order of preference):

- 1 By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
- 2 The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
- 3 The ITAL selected router-id.
- 4 The primary IPv4 address of an interface over which this OSPF process is running. The first interface address in the OSPF interface is selected.

**Note**

Unlike OSPF version 3, OSPF version 2 is guaranteed to have at least one interface with an IPv4 address configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to assign the IP address of 172.20.10.10 to the OSPF process 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
RP/0/RSP0/CPU0:router(config-ospf)# router-id 172.20.10.10
```

Related Commands

| Command | Description |
|---|---|
| clear ospf process, on page 946 | Resets an OSPF router process without stopping and restarting it. |
| ipv4 address | Sets a primary IPv4 address for an interface. |

router ospf

To configure an Open Shortest Path First (OSPF) routing process, use the **router ospf** command in global configuration mode. To terminate an OSPF routing process, use the **no** form of this command.

router ospf *process-name*

no router ospf *process-name*

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | Name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters without spaces. |
|---------------------|---|

Command Default

No OSPF routing process is defined.

Command Modes

global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can specify multiple OSPF routing processes in each router. Up to 10 processes can be configured. The recommendation is not to exceed 4 OSPF processes.

All OSPF configuration commands must be configured under an OSPF routing process. For example, two of these commands are the **default-metric** command and the **router-id** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |
| rib | read, write |

Examples

The following example shows how to instantiate an OSPF routing process called 109:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
```

Related Commands

| Command | Description |
|--|---|
| area (OSPF), on page 934 | Configures an OSPF area. |
| default-metric (OSPF), on page 966 | Sets default metric values for routes redistributed from another protocol into the OSPF protocol. |
| interface (OSPF), on page 996 | Defines the interfaces on which the OSPF protocol runs. |
| router-id (OSPF), on page 1076 | Configures a router ID for the OSPF process. |

security ttl (OSPF)

To set the security time-to-live (TTL) value in the IP header for Open Shortest Path First (OSPF) packets, use the **security ttl** command in the appropriate configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

security ttl [**hops** *hops-number*]

no security ttl

Syntax Description

| | |
|--------------------------------|--|
| hops <i>hops-number</i> | IP hops. Maximum number of hops allowed. Range is 1 to 254 hops. |
|--------------------------------|--|

Command Default

hops-number : 1

Command Modes

Router configuration
Area configuration
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **security ttl** command is used for the Generalized TTL Security Mechanism (GTSM) feature to prevent network attacks.

During the act of receiving Link State Advertisement (LSA) from neighbors, network attacks can occur because there are no checks that unicast or multicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the GTSM, RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to when OSPF packets are originated and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows how to set the security TTL for an interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet0/6/0/3
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# security ttl 2
```

sham-link

To configure an Open Shortest Path First OSPF sham link between two provider edge routers, use the **sham-link** command in VRF area configuration mode. To terminate an (OSPF) sham link, use the **no** form of this command.

sham-link *source-address destination-address*

no sham-link

Syntax Description

| | |
|----------------------------|--|
| <i>source-address</i> | IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation. |
| <i>destination-address</i> | IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation. |

Command Default

No sham link is configured.

Command Modes

VRF area configuration.

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure an OSPF sham link:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 109
```

```
RP/0/RSP0/CPU0:router(config_ospf)# vrf vrf_a
RP/0/RSP0/CPU0:router(config_ospf_vrf)# area 0
RP/0/RSP0/CPU0:router(config_ospf_vrf_ar)# sham-link 192.168.40.0 172.16.30.0
RP/0/RSP0/CPU0:router(config_ospf_vrf_ar_sl)# cost 23
```

Related Commands

| Command | Description |
|--|---|
| area (OSPF), on page 934 | Configures an OSPF area. |
| cost (OSPF), on page 954 | Explicitly specifies the cost of the interface (network) for OSPF path calculation. |
| vrf (OSPF), on page 1169 | Configures an OSPF VPN routing and forwarding (VRF) instance. |

show ospf

To display general information about Open Shortest Path First (OSPF) routing processes, use the **show ospf** command in EXEC mode.

show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] [**summary**]

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved values of the <i>vrf-name</i> argument. |
| all | (Optional) Specifies all OSPF VRF instances. |
| summary | (Optional) Displays OSPF summary information. |

Command Default

IPv4 and unicast address prefixes

Command Modes

EXEC
OSPFv3

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.3.0 | The show ospf vrf <i>vrf_name</i> command output was modified to display the VRF lite configuration status. |
| Release 5.1 | The output of the command was modified to include OSPFv3 status. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf** command to provide basic information about the OSPF processes running on the router. Additional options provide in-depth information.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf** command:

```
RP/0/RSP0/CPU0:router#show ospf

Routing Process "ospf 1" with ID 1.1.1.1
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
  Initial SPF schedule delay 5000 msec
  Minimum hold time between two consecutive SPF's 10000 msec
  Maximum wait time between two consecutive SPF's 10000 msec
  Initial LSA throttle delay 500 msec
  Minimum hold time for LSA throttle 5000 msec
  Maximum wait time for LSA throttle 5000 msec
  Minimum LSA interval 5000 msec. Minimum LSA arrival 1 sec
  Maximum number of configured interfaces 255
  Number of external LSA 0. Checksum Sum 00000000
  Number of opaque AS LSA 0. Checksum Sum 00000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  External flood list length 0
  Non-Stop Forwarding enabled
    Area BACKBONE(0) (Inactive)
      Number of interfaces in this area is 2
      SPF algorithm executed 8 times
      Number of LSA 2. Checksum Sum 0x01ba83
      Number of opaque link LSA 0. Checksum Sum 00000000
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 0
      Flood list length 0
    Area 1
      Number of interfaces in this area is 1
      SPF algorithm executed 9 times
      Number of LSA 2. Checksum Sum 0x0153ea
      Number of opaque link LSA 0. Checksum Sum 00000000
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 0
      Flood list length 0
```

This table describes the significant fields shown in the display.

Table 89: show ospf Field Descriptions

| Field | Description |
|---|---|
| Routing Process "ospf 201" with ID 172.22.110.200 | OSPF process name. |
| Supports only | Number of types of service supported (Type 0 only). |
| It is | Types are internal, area border, or autonomous system boundary. |
| Redistributing External Routes from | Lists of redistributed routes, by protocol. |
| SPF schedule delay | Delay time of SPF calculations. |
| Minimum LSA interval | Minimum interval between LSAs. |
| Minimum LSA arrival | Minimum elapsed time between accepting an update for the same link-state advertisement (LSA). |
| external LSA | Total number of Type 5 LSAs in the LSDB. |
| opaque LSA | Total number of Type 10 LSAs in the LSDB. |
| DCbitless...AS LSA | Total number of Demand Circuit Type 5 and Type 11 LSAs. |
| DoNotAge...AS LSA | Total number of Type 5 and Type 11 LSAs with the DoNotAge bit set. |
| Number of areas | Number of areas in router, area addresses, and so on. |
| Area BACKBONE | Backbone is area 0. |

This sample output from the **show ospf vrf vrf_name** command displays the VRF Lite status:

```
RP/0/RSP0/CPU0:router#show ospf vrf vrf1

VRF vrf1 in Routing Process "ospf 100" with ID 1.1.1.1
NSR (Non-stop routing) is Disabled
Supports only single TOS(TOS0) routes
Supports opaque LSA
It is an area border router
VRF Lite is enabled
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 50 msec
Minimum hold time between two consecutive SPFs 200 msec
Maximum wait time between two consecutive SPFs 5000 msec
Initial LSA throttle delay 50 msec
Minimum hold time for LSA throttle 200 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA interval 200 msec. Minimum LSA arrival 100 msec
LSA refresh interval 1800 seconds
Flood pacing interval 33 msec. Retransmission pacing interval 66 msec
Adjacency stagger enabled; initial (per area): 2, maximum: 64
  Number of neighbors forming: 0, 2 full
Maximum number of configured interfaces 1024
```

```
Number of external LSA 0. Checksum Sum 00000000
Number of opaque AS LSA 0. Checksum Sum 00000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
External flood list length 0
SNMP trap is disabled
  Area BACKBONE(0)
    Number of interfaces in this area is 1
    SPF algorithm executed 4 times
    Number of LSA 16. Checksum Sum 0x071c6a
    Number of opaque link LSA 0. Checksum Sum 00000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
    Number of LFA enabled interfaces 0, LFA revision 0
    Number of Per Prefix LFA enabled interfaces 0
    Number of neighbors forming in staggered mode 0, 1 full
  Area 1
    Number of interfaces in this area is 4
    SPF algorithm executed 5 times
    Number of LSA 14. Checksum Sum 0x066d93
    Number of opaque link LSA 0. Checksum Sum 00000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
    Number of LFA enabled interfaces 0, LFA revision 0
    Number of Per Prefix LFA enabled interfaces 0
    Number of neighbors forming in staggered mode 0, 1 full
```

show ospf border-routers

To display the internal Open Shortest Path First (OSPF) routing table entries to an Area Border Router (ABR) and Autonomous System Boundary Router (ASBR), use the **show ospf border-routers** command in EXEC mode.

show ospf [*process-name*] [**vrf** { *vrf-name* | **all** }] **border-routers** [*router-id*]

Syntax Description

| | |
|---------------------|--|
| <i>process-name</i> | (Optional) OSPF process name. If this argument is included, only information for the specified routing process is included. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved vrf-names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>router-id</i> | (Optional) Router ID associated with the border router. The value of the <i>router-id</i> argument can be any 32-bit router ID value specified in four-part, dotted-decimal notation. No default exists. |

Command Default

IPv4 and unicast address prefixes

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf border-routers** command to list all OSPF border routers visible to the specified processes and to ascertain the OSPF topology of the router.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf border-routers** command:

```
RP/0/RSP0/CPU0:router# show ospf border-routers

OSPF 1 Internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 172.31.97.53 [1] via 172.16.1.53, GigabitEthernet POS 3/0/0/0, ABR/ASBR , Area 0, SPF
3
```

This table describes the significant fields shown in the display.

Table 90: show ospf border-routers Field Descriptions

| Field | Description |
|-------------------------|---|
| i | Type of this route; i indicates an intra-area route, I an interarea route. |
| 172.31.97.53 | Router ID of destination. |
| [1] | Cost of using this route. |
| 172.16.1.53 | Next-Next hop toward the destination. |
| GigabitEthernet 3/0/0/0 | Packets destined for 172.16.1.53 are sent over GigabitEthernet interface 3/0/0/0. |
| ABR/ASBR | Router type of the destination; it is either an Area Border Router (ABR) or Autonomous System Boundary Router (ASBR) or both. |
| Area 0 | Area ID of the area from which this route is learned. |
| SPF 3 | Internal number of the shortest path first (SPF) calculation that installs this route. |

show ospf database

To display lists of information related to the Open Shortest Path First (OSPF) database for a specific router, use the **show ospf database** command in EXEC mode.

```

show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ adv-router ip-address ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ asbr-summary ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ asbr-summary ] [ link-state-id ]
[ internal ] [ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ asbr-summary ] [ link-state-id ]
[ internal ] [ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ database-summary ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ external ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ external ] [ link-state-id ] [ internal ]
[ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ external ] [ link-state-id ] [ internal ]
[ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ network ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ network ] [ link-state-id ] [ internal ]
[ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ network ] [ link-state-id ] [ internal ]
[ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ nssa-external ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ nssa-external ] [ link-state-id ] [ internal ]
[ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ nssa-external ] [ link-state-id ] [ internal ]
[ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-area ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-area ] [ link-state-id ] [ internal ]
[ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-area ] [ link-state-id ] [ internal ]
[ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-as ] [ link-state-id ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-as ] [ link-state-id ] [ internal ]
[ adv-router [ ip-address ] ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-as ] [ link-state-id ] [ internal ]
[ self-originate ]
show ospf [ process-name ] [ vrf { vrf-name | all } ] [ area-id ] database [ opaque-link ] [ link-state-id ]

```

```

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [opaque-link] [ link-state-id ] [internal]
[adv-router [ ip-address ]]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [opaque-link] [ link-state-id ] [internal]
[self-originate]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [router] [ link-state-id ]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [router] [internal] [adv-router
[ ip-address ]]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [router] [internal] [self-originate]
[ link-state-id ]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [self-originate]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [summary] [ link-state-id ]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [summary] [ link-state-id ] [internal]
[adv-router [ ip-address ]]

show ospf [ process-name ] [vrf {vrf-name|all}] [ area-id ] database [summary] [ link-state-id ] [internal]
[self-originate] [ link-state-id ]

```

Syntax Description

| | |
|-------------------------------------|--|
| <i>process-name</i> | (Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| adv-router <i>ip-address</i> | (Optional) Displays all LSAs of the specified router. |
| asbr-summary | (Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs. |

| | |
|-------------------------|---|
| <i>link-state-id</i> | <p>(Optional) Portion of the Internet environment that is being described by the advertisement. The value entered depends on the link-state type of the advertisement. It must be entered in the form of an IP address.</p> <p>When the link-state advertisement (LSA) is describing a network, the <i>link-state-id</i> can take one of two forms:</p> <ul style="list-style-type: none"> • The network IP address (as in Type 3 summary link advertisements and in autonomous system external link advertisements). • A derived address obtained from the link-state ID. <p>Note Masking the link-state ID of a network link advertisement with the subnet mask of the network yields the IP address of the network.</p> <p>When the LSA is describing a router, the link-state ID is always the OSPF router ID of the described router.</p> <p>When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link-state ID is set to Default Destination (0.0.0.0).</p> |
| internal | (Optional) Displays internal LSA information. |
| self-originate | (Optional) Displays only self-originated LSAs (from the local router). |
| database-summary | (Optional) Displays how many of each type of LSA for each area there are in the database and the total. |
| external | (Optional) Displays information only about the external LSAs. |
| network | (Optional) Displays information only about the network LSAs. |
| nssa-external | (Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs. |
| opaque-area | (Optional) Displays information about the opaque Type 10 LSAs. Type 10 denotes an area-local scope. Refer to RFC 2370 for more information on the opaque LSA options. |
| opaque-as | (Optional) Displays information about the opaque Type 11 LSAs. Type 11 denotes that the LSA is flooded throughout the autonomous system. |
| opaque-link | (Optional) Displays information about the opaque Type 9 LSAs. Type 9 denotes a link-local scope. |
| router | (Optional) Displays information only about the router LSAs. |
| summary | (Optional) Displays information only about the summary LSAs. |

Command Default

IPv4 and unicast address prefixes

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The various forms of the **show ospf database** command deliver information about different OSPF link-state advertisements. This command can be used to examine the link-state database (LSD) and its contents. Each router participating in an area having identical database entries pertaining to that area (with the exception of LSAs that are being flooded). Numerous options (such as **network** and **router**) are used to display portions of the database.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | ospf | read |

Examples The following is sample output from the **show ospf database** command when no arguments or keywords are used:

```
RP/0/RSP0/CPU0:router# show ospf database
OSPF Router with ID (172.20.1.11) (Process ID 1)
Router Link States (Area 0)
Link ID          ADV Router      Age             Seq#            Checksum Link count
172.20.1.8       172.20.1.8     1381           0x8000010D     0xEF60    2
172.20.1.11      172.20.1.11    1460           0x800002FE     0xEB3D    4
172.20.1.12      172.20.1.12    2027           0x80000090     0x875D    3
172.20.1.27      172.20.1.27    1323           0x800001D6     0x12CC    3
Net Link States (Area 0)
Link ID          ADV Router      Age             Seq#            Checksum
172.22.1.27     172.20.1.27    1323           0x8000005B     0xA8EE
172.22.1.11     172.20.1.11    1461           0x8000005B     0x7AC
Type-10 Opaque Link Area Link States (Area 0)
Link ID          ADV Router      Age             Seq#            Checksum Opaque ID
10.0.0.0         172.20.1.11    1461           0x800002C8     0x8483    0
10.0.0.0         172.20.1.12    2027           0x80000080     0xF858    0
10.0.0.0         172.20.1.27    1323           0x800001BC     0x919B    0
10.0.0.1         172.20.1.11    1461           0x8000005E     0x5B43    1
```

This table describes the significant fields shown in the display.

Table 91: show ospf database Field Descriptions

| Field | Description |
|------------|---|
| Link ID | Router ID number. |
| ADV Router | ID of the advertising router. |
| Age | Link-state age. |
| Seq# | Link-state sequence number (detects old or duplicate LSAs). |
| Checksum | Fletcher checksum of the complete contents of the LSA. |
| Link count | Number of interfaces detected for the router. |
| Opaque ID | Opaque LSA ID number. |

The following is sample output from the **show ospf database** command with the **asbr-summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database asbr-summary
OSPF Router with ID (192.168.0.1) (Process ID 300)
Summary ASB Link States (Area 0.0.0.0)
  LS age: 1463
  Options: (No TOS-capability)
  LS Type: Summary Links (AS Boundary Router)
  Link State ID: 172.17.245.1 (AS Boundary Router address)
  Advertising Router: 172.17.241.5
  LS Seq Number: 80000072
  Checksum: 0x3548
  Length: 28
  Network Mask: /0
  TOS: 0 Metric: 1
```

This table describes the significant fields shown in the display.

Table 92: show ospf database asbr-summary Field Descriptions

| Field | Description |
|---------------------|--|
| OSPF Router with ID | Router ID number. |
| Process ID | OSPF process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only). |
| LS Type | Link-state type. |

| Field | Description |
|--------------------|--|
| Link State ID | Link-state ID (ASBR). |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSAs. |
| Network Mask | Network mask implemented. |
| TOS | Type of service. |
| Metric | Link-state metric. |

The following is sample output from the **show ospf database** command with the **external** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database external
OSPF Router with ID (192.168.0.1) (Process ID 300)
      Type-5 AS External Link States

LS age: 280
Options: (No TOS-capability)
LS Type: AS External Link
Link State ID: 172.17.0.0 (External Network Number)
Advertising Router: 172.17.70.6
LS Seq Number: 80000AFD
Checksum: 0xC3A
Length: 36
Network Mask: 255.255.0.0
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 1
      Forward Address: 0.0.0.0
      External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 93: show ospf database external Field Descriptions

| Field | Description |
|----------------------------|--|
| OSPF Router with Router ID | Router ID number. |
| Process ID | OSPF process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only). |

| Field | Description |
|--------------------|--|
| LS Type | Link-state type. |
| Link State ID | Link-state ID (external network number). |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence number (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Network Mask | Network mask implemented. |
| Metric Type | External type. |
| TOS | Type of service. |
| Metric | Link-state metric. |
| Forward Address | Forwarding address. Data traffic for the advertised destination is forwarded to this address. If the forwarding address is set to 0.0.0.0, data traffic is forwarded instead to the originator of the advertisement. |
| External Route Tag | External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPF protocol itself. |

The following is sample output from the **show ospf database** command with the **network** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database network
OSPF Router with ID (192.168.0.1) (Process ID 300)
Net Link States (Area 0.0.0.0)
LS age: 1367
Options: (No TOS-capability)
LS Type: Network Links
Link State ID: 172.23.1.3 (address of Designated Router)
Advertising Router: 192.168.0.1
LS Seq Number: 800000E7
Checksum: 0x1229
Length: 52
Network Mask: /24
    Attached Router: 192.168.0.1
    Attached Router: 172.23.241.5
    Attached Router: 172.23.1.1
    Attached Router: 172.23.54.5
```

Attached Router: 172.23.1.5

This table describes the significant fields shown in the display.

Table 94: show ospf database network Field Descriptions

| Field | Description |
|---------------------|--|
| OSPF Router with ID | Router ID number. |
| Process ID | OSPF process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only). |
| LS Type | Link-state type. |
| Link State ID | Link-state ID of the designated router. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence number (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Network Mask | Network mask implemented. |
| Attached Router | List of routers attached to the network, by IP address. |

The following is sample output, carrying Multiprotocol Label Switching traffic engineering (MPLS TE) specification information, from the **show ospf database** command with the **opaque-area** keyword and a *link-state-id* of adv-router:

```
RP/0/RSP0/CPU0:router# show ospf database opaque-area adv-router 172.20.1.12
  OSPF Router with ID (172.20.1.11) (Process ID 1)
      Type-10 Opaque Link Area Link States (Area 0)
          LS age: 224
          Options: (No TOS-capability, DC)
          LS Type: Opaque Area Link
          Link State ID: 1.0.0.0
          Opaque Type: 1
          Opaque ID: 0
          Advertising Router: 172.20.1.12
          LS Seq Number: 80000081
          Checksum: 0xF659
          Length: 132
          Fragment number : 0
```

```

MPLS TE router ID : 172.20.1.12

Link connected to Point-to-Point network
  Link ID : 172.20.1.11
  Interface Address : 172.21.1.12
  Neighbor Address : 172.21.1.11
  Admin Metric : 10
  Maximum bandwidth : 193000
  Maximum reservable bandwidth : 125000
  Number of Priority : 8
  Priority 0 : 125000      Priority 1 : 125000
  Priority 2 : 125000      Priority 3 : 125000
  Priority 4 : 125000      Priority 5 : 125000
  Priority 6 : 125000      Priority 7 : 100000
  Affinity Bit : 0x0

Number of Links : 1

```

The following is sample output from the **show ospf database** command that displays a Type 10, Router Information LSA:

```

RP/0/RSP0/CPU0:router# show ospf database opaque-area 4.0.0.0

      OSPF Router with ID (3.3.3.3) (Process ID orange)

      Type-10 Opaque Link Area Link States (Area 0)

LS age: 105
Options: (No TOS-capability, DC)
LS Type: Opaque Area Link
Link State ID: 4.0.0.0
Opaque Type: 4
Opaque ID: 0
Advertising Router: 3.3.3.3
LS Seq Number: 80000052
Checksum: 0x34e2
Length: 52
Fragment number: 0

Router Information TLV: Length: 4
Capabilities:
  Graceful Restart Helper Capable
  Traffic Engineering enabled area
  All capability bits: 0x50000000

PCE Discovery TLV: Length: 20
IPv4 Address: 3.3.3.3
PCE Scope: 0x20000000
Compute Capabilities:
  Inter-area default (Rd-bit)
  Compute Preferences:
  Intra-area: 0  Inter-area: 0
  Inter-AS: 0  Inter-layer: 0

```

This table describes the significant fields shown in the display.

Table 95: show ospf database opaque-area Field Descriptions

| Field | Description |
|---------------------|--------------------|
| OSPF Router with ID | Router ID number. |
| Process ID | OSPF process name. |

| Field | Description |
|------------------------------|---|
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only). |
| LS Type | Link-state type. |
| Link State ID | Link-state ID. |
| Opaque Type | Opaque link-state type. |
| Opaque ID | Opaque ID number. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Fragment number | Arbitrary value used to maintain multiple traffic engineering LSAs. |
| Link ID | Link ID number. |
| Interface Address | ID address of the interface. |
| Neighbor Address | IP address of the neighbor. |
| Admin Metric | Administrative metric value used by MPLS TE. |
| Maximum bandwidth | Specifies maximum bandwidth (in kbps). |
| Maximum reservable bandwidth | Specifies maximum reservable bandwidth (in kbps). |
| Number of Priority | Priority number. |
| Affinity Bit | Used by MPLS TE. |
| Router Information TLV | Router capabilities are advertised in this TLV. |
| Capabilities | Some router capabilities include stub router, traffic engineering, graceful restart, and graceful restart helper. |
| PCE Discovery TLV | PCE address and capability information is advertised in this TLV. |

| Field | Description |
|-----------------------------|--|
| IPv4 Address | Configured PCE IPv4 address. |
| PCE Scope | Computation capabilities of the PCE. |
| Compute Capabilities | Compute capabilities and preferences of the PCE. |
| Inter-area default (RD-bit) | PCE compute capabilities such as intra-area, inter-area, inter-area default, inter-AS, inter-AS default and inter-layer. |
| Compute Preferences | Order or preference of path computation that includes intra-area, inter-area, inter-AS, and inter-layer preferences. |

The following is sample output from the **show ospf database** command with the **router** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database router
OSPF Router with ID (192.168.0.1) (Process ID 300)
Router Link States (Area 0.0.0.0)
  LS age: 1176
  Options: (No TOS-capability)
  LS Type: Router Links
  Link State ID: 172.23.21.6
  Advertising Router: 172.23.21.6
  LS Seq Number: 80002CF6
  Checksum: 0x73B7
  Length: 120
  AS Boundary Router
  Number of Links: 8
  Link connected to: another Router (point-to-point)
  (Link ID) Neighboring Router ID: 172.23.21.5
  (Link Data) Router Interface address: 172.23.21.6
  Number of TOS metrics: 0
  TOS 0 Metrics: 2
```

This table describes the significant fields shown in the display.

Table 96: show ospf database router Field Descriptions

| Field | Description |
|---------------------|--|
| OSPF Router with ID | Router ID number. |
| Process ID | OSPF process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only). |
| LS Type | Link-state type. |

| Field | Description |
|--------------------|--|
| Link State ID | Link-state ID. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| AS Boundary Router | Definition of router type. |
| Number of Links | Number of active links. |
| Link ID | Link type. |
| Link Data | Router interface address. |
| TOS | Type of service metric (Type 0 only). |

The following is sample output from **show ospf database** command with the **summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database summary
      OSPF Router with ID (192.168.0.1) (Process ID 300)

Summary Net Link States (Area 0.0.0.0)

LS age: 1401
Options: (No TOS-capability)
LS Type: Summary Links (Network)
Link State ID: 172.23.240.0 (Summary Network Number)
Advertising Router: 172.23.241.5
LS Seq Number: 80000072
Checksum: 0x84FF
Length: 28
Network Mask: /24
      TOS: 0 Metric: 1
```

This table describes the significant fields shown in the display.

Table 97: show ospf database summary Field Descriptions

| Field | Description |
|---------------------|--------------------|
| OSPF Router with ID | Router ID number. |
| Process ID | OSPF process name. |
| LS age | Link-state age. |

| Field | Description |
|--------------------|--|
| Options | Type of service options (Type 0 only). |
| LS Type | Link-state type. |
| Link State ID | Link-state ID (summary network number). |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Network Mask | Network mask implemented. |
| TOS | Type of service. |
| Metric | Link-state metric. |

The following is sample output from **show ospf database** command with the **database-summary** keyword:

```
RP/0/RSP0/CPU0:router# show ospf database database-summary
      OSPF Router with ID (172.19.65.21) (Process ID 1)

Area 0 database summary
  LSA Type    Count  Delete  Maxage
  Router      2      0       0
  Network     1      0       0
  Summary Net 2      0       0
  Summary ASBR 0      0       0
  Type-7 Ext  0      0       0
  Opaque Link 0      0       0
  Opaque Area 0      0       0
  Subtotal    5      0       0

Process 1 database summary
  LSA Type    Count  Delete  Maxage
  Router      2      0       0
  Network     1      0       0
  Summary Net 2      0       0
  Summary ASBR 0      0       0
  Type-7 Ext  0      0       0
  Opaque Link 0      0       0
  Opaque Area 0      0       0
  Type-5 Ext  2      0       0
  Opaque AS   0      0       0
  Total       7      0       0
```

This table describes the significant fields shown in the display.

Table 98: show ospf database database-summary Field Descriptions

| Field | Description |
|--------------|---|
| LSA Type | Link-state type. |
| Count | Number of advertisements in that area for each link-state type. |
| Delete | Number of LSAs that are marked "Deleted" in that area. |
| Maxage | Number of LSAs that are marked "Maxaged" in that area. |

show ospf flood-list

To display a list of Open Shortest Path First (OSPF) link-state advertisements (LSAs) waiting to be flooded over an interface, use the **show ospf flood-list** command in EXEC mode.

```
show ospf [ process-name ] [ vrf {vrf-name| all} ] [ area-id ] flood-list [type interface-path-id]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

All interfaces

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf flood-list** command to display LSAs in flood queue and queue length.

Flood list information is transient and normally the flood lists are empty.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf flood-list** command for interface GigabitEthernet 3/0/0/0:

```
RP/0/RSP0/CPU0:router# show ospf flood-list GigabitEthernet 3/0/0/0

Interface GigabitEthernet3/0/0/0, Queue length 20
Link state retransmission due in 12 msec
Displaying 6 entries from flood list:

Type  LS ID          ADV RTR          Seq NO          Age          Checksum
 5  10.2.195.0        200.0.0.163     0x80000009     0           0xFB61
 5  10.1.192.0        200.0.0.163     0x80000009     0           0x2938
 5  10.2.194.0        200.0.0.163     0x80000009     0           0x757
 5  10.1.193.0        200.0.0.163     0x80000009     0           0x1E42
 5  10.2.193.0        200.0.0.163     0x80000009     0           0x124D
 5  10.1.194.0        200.0.0.163     0x80000009     0           0x134C
```

This table describes the significant fields shown in the display.

Table 99: show ospf flood-list Field Descriptions

| Field | Description |
|----------------------------------|---|
| GigabitEthernet3/0/0/0 | Interface for which information is displayed. |
| Queue length | Number of LSAs waiting to be flooded. |
| Link state retransmission due in | Length of time (in milliseconds) before next link-state transmission. |
| Type | Type of LSA. |
| LS ID | Link-state ID of the LSA. |
| ADV RTR | IP address of the advertising router. |
| Seq NO | Sequence number of the LSA. |
| Age | Age of the LSA (in seconds). |
| Checksum | Checksum of the LSA. |

show ospf interface

To display Open Shortest Path First (OSPF) interface information, use the **show ospf interface** command in EXEC mode.

```
show ospf [ process-name ] [ vrf {vrf-name| all} ] [ area-id ] interface [brief] [type interface-path-id]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| brief | (Optional) Displays brief interface information. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

All interfaces

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read |

Examples

The following is sample output from the **show ospf interface** command:

```
RP/0/RSP0/CPU0:router# show ospf interface

GigabitEthernet0/2/0/1 is up, line protocol is up
  Internet Address 121.10.10.2/24, Area 2
  Process ID 1, Router ID 200.2.2.2, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Index 1/3, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 3, maximum is 10
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 101.3.3.3
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 1
    Multi-Area interface exist in area 1 Neighbor Count is 1
GigabitEthernet0/3/0/0 is up, line protocol is up
  Internet Address 145.10.10.2/16, Area 3
  Process ID 1, Router ID 200.2.2.2, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  BFD enabled, BFD interval 15 msec, BFD multiplier 3
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Index 1/5, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 3, maximum is 11
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 101.3.3.3
  Suppress hello for 0 neighbor(s)
  Message digest authentication enabled
  Youngest key id is 1
  Multi-area interface Count is 0
```

The following is sample output when a multiple area adjacency is configured on GigabitEthernet interface 0/2/0/1:

```
RP/0/RSP0/CPU0:router#show ospf 1 interface
  GigabitEthernet0
    /2
    /0
    /1GigabitEthernet0/2/0/1 is up, line protocol is up
      Internet Address 121.10.10.2/24, Area 1
      Process ID 1, Router ID 200.2.2.2, Network Type POINT_TO_POINT, Cost: 1
      Interface is multi-area adjacency
      LDP Sync Enabled, Sync Status: Achieved
      Transmit Delay is 1 sec, State POINT_TO_POINT,
      Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
      Non-Stop Forwarding (NSF) enabled
        Hello due in 00:00:09
      Index 2/4, flood queue length 0
      Next 0(0)/0(0)
      Last flood scan length is 3, maximum is 10
```

```

Last flood scan time is 0 msec, maximum is 1 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 101.3.3.3
  Suppress hello for 0 neighbor(s)
Multi-area interface Count is 0
Loopback0 is up, line protocol is up
  Internet Address 200.2.2.2/32, Area 1
  Process ID 1, Router ID 200.2.2.2, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host
RP/0/RSP0/CPU0:router#

```

This table describes the significant fields shown in the display.

Table 100: show ospf interface Field Descriptions

| Field | Description |
|----------------------------|---|
| GigabitEthernet | Status of the physical link. |
| line protocol | Operational status of the protocol. |
| Internet Address | Interface IP address, subnet mask, and area address. |
| Process ID | OSPF process ID, router ID, network type, and link-state cost. |
| Transmit Delay | Transmit delay, interface state, and router priority. |
| Timer intervals configured | Configuration of timer intervals. |
| Hello | Number of seconds until next hello packet is sent over this interface. |
| Index | Area and autonomous system flood indexes. |
| Next 0 (0) /0 (0) | Next area and autonomous system flood information, data pointer, and index. |
| Last flood scan length | Length of last flood scan. |
| Last flood scan time | Time (in milliseconds) of last flood scan. |
| Neighbor Count | Count of network neighbors and list of adjacent neighbors. |
| Suppress hello | Count of neighbors suppressing hello messages. |
| Multi-area interface | Multiple area interface information for the primary interface, such as count and area/neighbor locations. |

show ospf mpls traffic-eng

To display information about the links and fragments available on the local router for traffic engineering, use the **show ospf mpls traffic-eng** command in EXEC mode.

```
show ospf [ process-name ] [vrf {vrf-name| all}] [ area-id ] [type interface-path-id] mpls traffic-eng {link|fragment}
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) OSPF process name that uniquely identifies an OSPF routing process. The process name is any alphanumeric string no longer than 40 characters. If this argument is included, only information for the specified routing process is included. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| link | Provides detailed information about the links over which traffic engineering is supported on the local router. |
| fragment | Provides detailed information about the traffic engineering fragments on the local router. |

Command Default All links or fragments

Command Modes EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf mpls traffic-eng** command when the **link** keyword is specified:

```
RP/0/RSP0/CPU0:router# show ospf mpls traffic-eng link
OSPF Router with ID (10.10.10.10) (Process ID 1)
Area 0 has 2 MPLS TE links. Area instance is 67441.
Links in hash bucket 3.
Link is associated with fragment 1. Link instance is 67441
Link connected to Point-to-Point network
Link ID : 10.10.10.8
Interface Address : 10.10.10.2
Neighbor Address : 10.10.10.3
Admin Metric : 0
Maximum bandwidth : 19440000
Maximum global pool reservable bandwidth : 25000000
Maximum sub pool reservable bandwidth : 3125000
Number of Priority : 8
Global pool unreserved BW
Priority 0 : 25000000 Priority 1 : 25000000
Priority 2 : 25000000 Priority 3 : 25000000
Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
Priority 4 : 3125000 Priority 5 : 3125000
Priority 6 : 3125000 Priority 7 : 3125000
Affinity Bit : 0
Links in hash bucket 8.
Link is associated with fragment 0. Link instance is 67441
Link connected to Point-to-Point network
Link ID : 10.1.1.1
Interface Address : 10.10.25.4
Neighbor Address : 10.10.25.5
Admin Metric : 0
Maximum bandwidth : 19440000
Maximum global pool reservable bandwidth : 25000000
Maximum sub pool reservable bandwidth : 3125000
Number of Priority : 8
Global pool unreserved BW
Priority 0 : 25000000 Priority 1 : 25000000
Priority 2 : 25000000 Priority 3 : 25000000
Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
```

```

Priority 4 :    3125000  Priority 5 :    3125000
Priority 6 :    3125000  Priority 7 :    3125000
Affinity Bit : 0

```

This table describes the significant fields shown in the display.

Table 101: show ospf mpls traffic-eng link Field Descriptions

| Field | Description |
|--|---|
| Link ID | Link type. |
| Interface address | IP address of the interface. |
| Neighbor address | IP address of the neighbor. |
| Admin Metric | Administrative distance metric value used by Multiprotocol Label Switching traffic engineering (MPLS TE). |
| Maximum bandwidth | Bandwidth capacity of the link (in kbps). |
| Maximum global pool reservable bandwidth | Maximum amount of bandwidth that is available for reservation in the global pool. |
| Maximum sub pool reservable bandwidth | Maximum amount of bandwidth that is available for reservation in the subpool. |
| Number of Priority | Priority number. |
| Global pool unreserved BW | Amount of unreserved bandwidth that is available in the global pool. |
| Sub pool unreserved BW | Amount of unreserved bandwidth that is available in the subpool. |
| Affinity Bit | Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1. |

The following is sample output from the **show ospf mpls traffic-eng** command when the **fragment** keyword is specified:

```

RP/0/RSP0/CPU0:router# show ospf mpls traffic-eng fragment

      OSPF Router with ID (10.10.10.10) (Process ID 1)

Area 0 has 2 MPLS TE fragment. Area instance is 67441.
MPLS router address is 10.10.10.10
Next fragment ID is 2

```

```

Fragment 0 has 1 link. Fragment instance is 67441.
Fragment has 1 link the same as last update.
Fragment advertise MPLS router address
Link is associated with fragment 0. Link instance is 67441
Link connected to Point-to-Point network
Link ID : 10.1.1.1
Interface Address : 10.10.25.4
Neighbor Address : 10.10.25.5
Admin Metric : 0
Maximum bandwidth : 19440000
Maximum global pool reservable bandwidth : 25000000
Maximum sub pool reservable bandwidth : 3125000
Number of Priority : 8
Global pool unreserved BW
Priority 0 : 25000000 Priority 1 : 25000000
Priority 2 : 25000000 Priority 3 : 25000000
Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
Priority 4 : 3125000 Priority 5 : 3125000
Priority 6 : 3125000 Priority 7 : 3125000
Affinity Bit : 0

```

```

Fragment 1 has 1 link. Fragment instance is 67441.
Fragment has 0 link the same as last update.
Link is associated with fragment 1. Link instance is 67441
Link connected to Point-to-Point network
Link ID : 10.10.10.8
Interface Address : 10.10.10.2
Neighbor Address : 10.10.10.3
Admin Metric : 0
Maximum bandwidth : 19440000
Maximum global pool reservable bandwidth : 25000000
Maximum sub pool reservable bandwidth : 3125000
Number of Priority : 8
Global pool unreserved BW
Priority 0 : 25000000 Priority 1 : 25000000
Priority 2 : 25000000 Priority 3 : 25000000
Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
Priority 4 : 3125000 Priority 5 : 3125000
Priority 6 : 3125000 Priority 7 : 3125000
Affinity Bit : 0

```

This table describes the significant fields shown in the display.

Table 102: show ospf mpls traffic-eng fragment Field Descriptions

| Field | Description |
|-------------------|--|
| Area instance | Number of times traffic engineering information or any link changed. |
| Link instance | Number of times any link changed. |
| Link ID | Link type. |
| Interface address | IP address of the interface. |

| Field | Description |
|--|---|
| Neighbor address | IP address of the neighbor. |
| Admin Metric | Administrative distance metric value used by MPLS TE. |
| Maximum bandwidth | Bandwidth capacity of the link (in kbps). |
| Maximum global pool reservable bandwidth | Maximum amount of bandwidth that is available for reservation in the global pool. |
| Maximum sub pool reservable bandwidth | Maximum amount of bandwidth that is available for reservation in the subpool. |
| Number of Priority | Priority number. |
| Global pool unreserved BW | Amount of unreserved bandwidth that is available in the global pool. |
| Sub pool unreserved BW | Amount of unreserved bandwidth that is available in the subpool. |
| Affinity Bit | Used by MPLS TE. Attribute values required for links carrying this tunnel. A 32-bit dotted-decimal number. Valid values are from 0x0 to 0xFFFFFFFF, representing 32 attributes (bits), where the value of an attribute is 0 or 1. |

show ospf message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospf message-queue** command in EXEC mode.

show ospf message-queue

This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.9.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | ospf | read |

Examples The following is sample output from the **show ospf message-queue** command:

```
RP/0/RSP0/CPU0:router# show ospf 1 message-queue

OSPF 1
  Hello Input Queue:
    Current queue length: 0
    Event scheduled: 0
    Total queuing failures: 0
    Maximum length : 102
    Pkts pending processing: 0
    Limit: 5000

  Router Message Queue
    Current instance queue length: 0
    Current redistribution queue length: 0
    Current ex spf queue length: 0
    Current sum spf queue length: 0
    Current intra spf queue length: 0
    Event scheduled: 0
    Maximum length : 101
    Total low queuing failures: 0
```

```

Total medium queuing failures: 0
Total high queuing failures: 0
Total instance events: 919
Processing quantum : 300
Low queuing limit: 8000
Medium queuing limit: 9000
High queuing limit: 9500
Rate-limited LSA processing quantum: 150
Current rate-limited LSA queue length: 0
Rate-limited LSA queue peak len: 517

Rate-limited LSAs processed: 4464
Flush LSA processing quantum: 150
Current flush LSA queue length: 0
Flush LSA queue peak len: 274
Rate-limited flush LSAs processed: 420

SPF-LSA-limit processing quantum: 150
Managed timers processing quantum: 50
Instance message count: 0
Instance pulse send count: 919
Instance pulse received count: 919
Global pulse count: 0
Instance Pulse errors: 0

TE Message Queue
Current queue length: 0
Total queuing failures: 0
Maximum length : 0

Number of Dlink errors: 0

```

This table describes the significant fields shown in the display.

Table 103: show ospf message-queue Field Descriptions

| Field | Description |
|------------------------|---|
| Hello Input Queue | This section provides statistics on the number of events and incoming packets processed in the Hello (incoming packet) thread of the OSPF process. |
| Router Message Queue | This section provides statistics on the events and messages processed in the Router (primary) thread of the OSPF process. |
| TE Message Queue | This section provides statistics on traffic-engineering events and messages received by OSPF from TE (the te_control process). These events are processed in the Router thread of the OSPF process. |
| Number of Dlink errors | The number of enqueueing or dequeuing errors seen across all the linked-lists in the OSPF process. |

Related Commands

| Command | Description |
|---|---|
| queue dispatch incoming, on page 1057 | Limits the number of continuous incoming events processed. |
| queue dispatch rate-limited-lsa, on page 1059 | Sets the maximum number of rate-limited link-state advertisements (LSAs) processed per run. |
| queue dispatch spf-lsa-limit, on page 1061 | Limits the number of summary or external Type 3 to Type 7 link-state advertisements (LSAs) processed per shortest path first (SPF) run. |
| queue limit, on page 1063 | Sets the high watermark for incoming priority events. |

show ospf neighbor

To display Open Shortest Path First (OSPF) neighbor information on an individual interface basis, use the **show ospf neighbor** command in EXEC mode.

```
show ospf [ process-name ] [ vrf {vrf-name| all} ] [ area-id ] neighbor [[type interface-path-id] [ neighbor-id ]
[detail]] area-sorted]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) Neighbor ID. |
| detail | (Optional) Displays all neighbors given in detail (lists all neighbors). |
| area-sorted | (Optional) Specifies that all neighbors are grouped by area. |

Command Default

All neighbors

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf neighbor** command showing two lines of summary information for each neighbor:

```
RP/0/RSP0/CPU0:router# show ospf neighbor
Neighbors for OSPF
Neighbor ID      Pri   State           Dead Time   Address           Interface
192.168.199.137 1     FULL/DR         0:00:31    172.31.80.37     GigabitEthernet 0/3/0/2
Neighbor is up for 18:45:22
192.168.48.1    1     FULL/DROTHER    0:00:33    192.168.48.1     GigabitEthernet 0/3/0/3
Neighbor is up for 18:45:30
192.168.48.200 1     FULL/DROTHER    0:00:33    192.168.48.200   GigabitEthernet 0/3/0/3
Neighbor is up for 18:45:25
192.168.199.137 5     FULL/DR         0:00:33    192.168.48.189   GigabitEthernet 0/3/0/3
Neighbor is up for 18:45:27
```

This table describes the significant fields shown in the display.

Table 104: show ospf neighbor Field Descriptions

| Field | Description |
|----------------|---|
| Neighbor ID | Neighbor router ID. |
| Pri | Designated router priority. |
| State | OSPF state. |
| Dead time | Time (in hh:mm:ss) that must elapse before OSPF declares the neighbor dead. |
| Address | Address of next hop. |
| Interface | Interface name of next hop. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |

The following is sample output showing summary information about the neighbor that matches the neighbor ID:

```
RP/0/RSP0/CPU0:router# show ospf neighbor 192.168.199.137

Neighbor 192.168.199.137, interface address 172.31.80.37
  In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x2
  Dead timer due in 0:00:32
  Neighbor is up for 18:45:30
  Number of DBD retrans during last exchange 0
  Index 1/1, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum 0 msec
Neighbor 192.168.199.137, interface address 192.168.48.189
  In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/3
  Neighbor priority is 5, State is FULL, 6 state changes
  Options is 0x2
  Dead timer due in 0:00:32
  Neighbor is up for 18:45:30
  Number of DBD retrans during last exchange 0
  Index 1/1, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 105: show ospf neighbor 192.168.199.137 Field Descriptions

| Field | Description |
|-------------------|---|
| Neighbor | Neighbor router ID. |
| interface address | IP address of the interface. |
| In the area | Area and interface through which the OSPF neighbor is known. |
| Neighbor priority | Router priority of neighbor and neighbor state. |
| State | OSPF state. |
| state changes | Number of state changes for this neighbor. |
| DR is | Neighbor ID of the designated router. |
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents(E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub.) |

| Field | Description |
|-----------------------|--|
| Dead timer | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |
| Number of DBD retrans | Number of re-sent database description packets. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

If you specify the interface along with the neighbor ID, the software displays the neighbors that match the neighbor ID on the interface, as in the following sample display:

```
RP/0/RSP0/CPU0:router# show ospf neighbor GigabitEthernet 0/3/0/2 192.168.199.137

Neighbor 192.168.199.137, interface address 172.31.80.37
  In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x2
  Dead timer due in 0:00:32
  Neighbor is up for 18:45:30
  Number of DBD retrans during last exchange 0
  Index 1/1, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 106: show ospf neighbor GigabitEthernet 0/3/0/2 192.168.199.137 Field Descriptions

| Field | Description |
|-------------------|--|
| Neighbor | Neighbor router ID. |
| interface address | IP address of the interface. |
| In the area | Area and interface through which the OSPF neighbor is known. |
| Neighbor priority | Router priority of the neighbor. |
| State | OSPF state. |
| state changes | Number of state changes for this neighbor. |
| DR is | Neighbor ID of the designated router. |

| Field | Description |
|-----------------------|---|
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub) |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |
| Number of DBD retrans | Number of re-sent database description packets. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

You can also specify the interface without the neighbor ID to show all neighbors on the specified interface, as in the following sample display:

```
RP/0/RSP0/CPU0:router# show ospf neighbor GigabitEthernet POS 0/3/0/3

Neighbors for OSPF ospf1

   ID          Pri   State          Dead Time    Address      Interface
192.168.48.1   1    FULL/DROTHER  0:00:33     192.168.48.1 GigabitEthernet POS
0/3/0/3
   Neighbor is up for 18:50:52
192.168.48.200 1    FULL/DROTHER  0:00:32     192.168.48.200 GigabitEthernet POS
0/3/0/3
   Neighbor is up for 18:50:52
192.168.199.137 5    FULL/DR       0:00:32     192.168.48.189 GigabitEthernet POS
0/3/0/3
   Neighbor is up for 18:50:52

Total neighbor count: 3
```

This table describes the significant fields shown in the display.

Table 107: show ospf neighbor GigabitEthernet 0/3/0/3 Field Descriptions

| Field | Description |
|-----------|--|
| ID | Neighbor router ID. |
| Pri | Route priority of the neighbor. |
| State | OSPF state. |
| Dead Time | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |

| Field | Description |
|-----------------------|---|
| Address | Address of next hop. |
| Interface | Interface name of next hop. |
| Neighbor is up | Time (in hh:mm:ss) that the OSPF neighbor has been up. |
| Options | Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates area is not a stub; 0 indicates area is a stub) |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |
| Number of DBD retrans | Number of re-sent database description packets. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

The following samples are from output from the **show ospf neighbor detail** command:

```
RP/0/RSP0/CPU0:router# show ospf neighbor detail

Neighbor 192.168.199.137, interface address 172.31.80.37
  In the area 0.0.0.0 via interface GigabitEthernet 0/3/0/2
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x2
  Dead timer due in 0:00:32
  Neighbor is up for 18:45:30
  Number of DBD retrans during last exchange 0
  Index 1/1, retransmission queue length 0, number of retransmission 0
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum 0 msec

Total neighbor count: 1

Neighbor 10.1.1.1, interface address 192.168.13.1
  In the area 0 via interface GigabitEthernet0/3/0/1
  Neighbor priority is 1, State is FULL, 10 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x52
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:36
  Neighbor is up for 1w2d
  Number of DBD retrans during last exchange 0
  Index 3/3, retransmission queue length 0, number of retransmission 5
  First 0(0)/0(0) Next 0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

```

Neighbor 10.4.4.4, interface address 192.168.34.4
  In the area 0 via interface GigabitEthernet0/3/0/2
  Neighbor priority is 1, State is FULL, 48 state changes
  DR is 0.0.0.0 BDR is 0.0.0.0
  Options is 0x12
  LLS Options is 0x1 (LR)
  Dead timer due in 00:00:30
  Neighbor is up for 00:40:03
  Number of DBD retrans during last exchange 0
  Index 2/2, retransmission queue length 0, number of retransmission 6
  First 0(0)/0(0) Next 0(0)/0(0)
  Last retransmission scan length is 0, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

```

This table describes the significant fields shown in the display.

Table 108: show ospf neighbor detail Field Descriptions

| Field | Description |
|-------------------------|---|
| Neighbor | Neighbor router ID. |
| interface address | IP address of the interface. |
| In the area | Area and interface through which the OSPF neighbor is known. |
| Neighbor priority | Router priority of neighbor and neighbor state. |
| State | OSPF state. |
| state changes | Number of state changes for this neighbor. |
| DR is | Neighbor ID of the designated router. |
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents. (E-bit only; possible values are 0 and 2; 2 indicates that the area is not a stub; 0 indicates that the area is a stub.) |
| LLS Options is 0x1 (LR) | Neighbor is NFS Cisco capable. |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |
| Number of DBD retrans | Number of re-sent database description packets. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

Related Commands

| Command | Description |
|--|-------------------------------------|
| router ospf , on page 1078 | Configures an OSPF routing process. |

show ospf request-list

To display the first ten link-state requests pending that the local router is making to the specified Open Shortest Path First (OSPF) neighbor and interface, use the **show ospf request-list** command in EXEC mode.

```
show ospf [ process-name ] [ vrf {vrf-name| all} ] [ area-id ] request-list [type interface-path-id] [ neighbor-id ]
```

Syntax Description

| | |
|---------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>i nterface-path-id</i> | Physical interface or virtual interface. Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPF neighbor. |

Command Default All neighbors

Command Modes EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency does not form between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that are suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Request list information is transient and normally the lists are empty.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf request-list** command:

```
RP/0/RSP0/CPU0:router# show ospf request-list 10.0.124.4 GigabitEthernet3/0/0/0
Request Lists for OSPF pagent
Neighbor 10.0.124.4, interface GigabitEthernet3/0/0/0 address 10.3.1.2
Type  LS ID          ADV RTR          Seq NO          Age  Checksum
  1   192.168.58.17     192.168.58.17   0x80000012     12  0x0036f3
  2   192.168.58.68     192.168.58.17   0x80000012     12  0x00083f
```

This table describes the significant fields shown in the display.

Table 109: show ospf request-list 10.0.124.4 GigabitEthernet3/0/0/0 Field Descriptions

| Field | Description |
|-----------|---|
| Neighbor | Specific neighbor receiving the request list from the local router. |
| Interface | Specific interface over which the request list is being sent. |
| Address | Address of the interface over which the request list is being sent. |
| Type | Type of link-state advertisement (LSA). |

| Field | Description |
|----------|---------------------------------------|
| LS ID | Link-state ID of the LSA. |
| ADV RTR | IP address of the advertising router. |
| Seq NO | Sequence number of the LSA. |
| Age | Age of the LSA (in seconds). |
| Checksum | Checksum of the LSA. |

Related Commands

| Command | Description |
|--|--|
| router ospf , on page 1078 | Configures an OSPF routing process. |
| show ospf retransmission-list , on page 1128 | Displays the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface. |

show ospf retransmission-list

To display the first ten link-state entries in the Open Shortest Path First (OSPF) retransmission list that the local router sends to the specified neighbor over the specified interface, use the **show ospf retransmission-list** command in EXEC mode.

```
show ospf [ process-name ] [ vrf {vrf-name| all} ] [ area-id ] retransmission-list [type interface-path-id]
[ neighbor-id ]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPF neighbor. |

Command Default

All neighbors

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Retransmission list information is transient, and normally the lists are empty.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf retransmission-list** command:

```
RP/0/RSP0/CPU0:router# show ospf retransmission-list 10.0.124.4 GigabitEthernet3/0/0/0
Neighbor 10.0.124.4, interface GigabitEthernet3/0/0/0 address 10.3.1.2
```

This table describes the significant fields shown in the display.

Table 110: show ospf retransmission-list 10.0.124.4 GigabitEthernet3/0/0/0 Field Descriptions

| Field | Description |
|-----------|---|
| Neighbor | Specified neighbor receiving the retransmission list from the local router. |
| Interface | Specified interface over which the retransmission list is being sent. |
| Address | Address of the interface. |

Related Commands

| Command | Description |
|--|-------------------------------------|
| router ospf , on page 1078 | Configures an OSPF routing process. |

| Command | Description |
|---|---|
| show ospf request-list , on page 1125 | Displays the first ten link-state requests pending that the local router is making to the specified neighbor and interface. |

show ospf routes

To display the Open Shortest Path First (OSPF) topology table, use the **show ospf routes** command in EXEC mode.

```
show ospf [ process-name ] [vrf {vrf-name| all}] routes [connected| external| local] [prefix mask]
[prefix/length] [multicast-intact]
```

Syntax Description

| | |
|-------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |
| connected | (Optional) Displays connected routes. |
| external | (Optional) Displays routes redistributed from other protocols. |
| local | (Optional) Displays the local routes redistributed from the Routing Information Base (RIB). |
| <i>prefix</i> | (Optional) IP prefix, which limits output to a specific route. If the <i>prefix</i> argument is specified, either the <i>length</i> or <i>mask</i> argument is required. |
| <i>mask</i> | (Optional) IP address mask. |
| <i>/ length</i> | (Optional) Prefix length, which can be indicated as a slash (/) and number. For example, /8 indicates that the first eight bits in the IP prefix are network bits. If <i>length</i> is used, the slash is required. |
| multicast-intact | (Optional) Displays multicast intact paths. |

Command Default

All route types

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf routes** command to display the OSPF private routing table (which contains only routes calculated by OSPF). If there is something wrong with a route in the RIB, then it is useful to check the OSPF copy of the route to determine if it matches the RIB contents. If it does not match, there is a synchronization problem between OSPF and the RIB. If the routes match and the route is incorrect, OSPF has made an error in its routing calculation.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf routes** command:

```
RP/0/RSP0/CPU0:router# show ospf routes
Topology Table for ospf 1 with ID 10.3.4.2
Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2
O E2 10.3.1.0/24, metric 1
     10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O   10.3.4.0/24, metric 1562
     10.3.4.2, directly connected, via GigabitEthernet 0/1/0/1
O E2 10.1.0.0/16, metric 1
     10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O IA 10.10.10.0/24, metric 1572
     10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
O E2 130.10.10.0/24, metric 20
     10.3.4.1, from 172.16.10.1, via GigabitEthernet 0/1/0/1
```

This table describes the significant fields shown in the display.

Table 111: show ospf route Field Descriptions

| Field | Description |
|-------|-----------------------------|
| O | OSPF route. |
| E | External Type 1 or 2 route. |

| Field | Description |
|-----------------------------|--|
| N | NSSA Type 1 or 2 |
| 10.3.1.0/24 | Network and subnet mask to which the local router has a route. |
| metric | Cost to reach network 10.3.1.0. |
| 10.3.4.1 | Next-hop router on the path to network 10.3.1.0. |
| from 172.16.10.1 | Router ID 172.16.10.1 is the router that advertised this route. |
| via GigabitEthernet 0/1/0/1 | Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/1/0/1. |

The following is sample output from the **show ospf routes** command with a process name of 100:

```
RP/0/RSP0/CPU0:router# show ospf 100 routes
Topology Table for ospf 100 with ID 172.23.54.14
Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2
O      10.1.5.0/24, metric 1562
       10.1.5.14, directly connected, via GigabitEthernet 0/3/0/3
O IA  21.0.0.0/24, metric 1572
       10.1.5.12, from 172.23.54.12, via GigabitEthernet 0/3/0/3
O      10.0.0.0/24, metric 10
       10.0.0.12, directly connected, via GigabitEthernet 0/2/0/3
```

This table describes the significant fields shown in the display.

Table 112: show ospf 100 route Field Descriptions

| Field | Description |
|-------------------|--|
| O | OSPF route. |
| IA | Interarea route. |
| 10.1.5.0/24 | Network and subnet mask to which the local router has a route. |
| metric 1562 | Cost to reach network 10.1.5.0. |
| 10.1.5.14 | Next-hop router on the path to network 10.1.5.0. |
| from 172.23.54.12 | Router ID 172.23.54.12 is the router that advertised this route. |

| Field | Description |
|-----------------------------|--|
| via GigabitEthernet 0/3/0/3 | Packets destined for the given prefix (10.3.1.0/24) are sent over GigabitEthernet interface 0/3/0/3. |

The following is sample output from the **show ospf routes** command with a prefix of 10.0.0.0 and a length of 24:

```
RP/0/RSP0/CPU0:router# show ospf routes 10.0.0.0/24
Topology Table for ospf 100 with ID 172.23.54.14
Codes:O - Intra area, O IA - Inter area
       O E1 - External type 1, O E2 - External type 2
       O N1 - NSSA external type 1, O N2 - NSSA external type 2
O IA 10.0.0.0/24, metric 1572
     10.1.5.12, from 172.23.54.12, via GigabitEthernet 0/3/0/3
```

This table describes the significant fields shown in the display.

Table 113: show ospf route 10.0.0.0/24 Field Descriptions

| Field | Description |
|-----------------------------|--|
| O | Route is an OSPF route. |
| IA | Route to network 10.0.0.0 is an interarea route. |
| 10.0.0.0/24 | Network and subnet mask to which the local router has a route. |
| metric 1572 | Cost to reach network 10.0.0.0. |
| 10.1.5.12 | IP address of next-hop router on the path to network 10.0.0.0. |
| from 172.23.54.12 | Router ID 172.23.54.12 is the router that advertised this route. |
| via GigabitEthernet 0/3/0/3 | Packets destined for the given prefix (10.0.0.0/24) are sent over GigabitEthernet interface 0/3/0/3. |

Related Commands

| Command | Description |
|--|--|
| router ospf , on page 1078 | Configures an OSPF routing process. |
| show route , on page 1412 | Displays current routes information in the Routing Information Base (RIB). |

| Command | Description |
|---|---|
| show rib opaques , on page 1394 | Displays opaque data installed in the Routing Information Base (RIB). |

show ospf sham-links

To display Open Shortest Path First (OSPF) sham-link information, use the **show ospf sham-links** command in EXEC mode.

show ospf [*process-name*] [**vrf** {*vrf-name*|**all**}] **sham-links**

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf sham-links** command to display OSPF sham-link information.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf sham-links** command:

```
RP/0/RSP0/CPU0:router# show ospf 1 vrf vrf_1 sham-links
```

```

Sham Links for OSPF 1, VRF vrf_1

Sham Link OSPF_SL0 to address 10.0.0.3 is up
Area 0, source address 10.0.0.1
IfIndex = 185
Run as demand circuit
DoNotAge LSA allowed., Cost of using 1
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Number of DBD retrans during last exchange 0
Index 2/2, retransmission queue length 0, number of retransmission 0
First 0(0)/0(0) Next 0(0)/0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
Keychain-based authentication enabled
Key id used is 2

```

This table describes the significant fields shown in the display.

Table 114: show ospf sham-links Field Descriptions

| Field | Description |
|--|--|
| Sham Link OSPF_SL0 to address | Address of the destination endpoint of the sham link. |
| IfIndex | ifindex associated with the sham link. |
| Run as demand circuit | Sham link is treated as a demand circuit. |
| DoNotAge LSA allowed | DoNotAge LSAs are allowed to be flooded over the sham link. |
| Cost of using | Sham-link cost. |
| Transmit Delay | Sham-link transmit delay. |
| State | Sham-link interface state. |
| Timer intervals configured | Various sham-link interface-related timers. |
| Hello due in | Time before the next Hello is sent over the sham link. |
| Adjacency State | State of the adjacency with the neighbor over the sham link. |
| Number of DBD retrans during last exchange | Number of DBD retransmissions during the last exchange over the sham link. |
| Index | Area flood index. |
| retransmission queue length | Retransmission queue length on the sham link. |

| Field | Description |
|---------------------------------------|--|
| number of retransmission | Number of retransmissions over the sham-link interface. |
| First | First flood information. |
| Next | Next flood information. |
| Last retransmission scan length is | Last retransmission scan length on the sham-link interface. |
| maximum is | Maximum retransmission scan length on the sham-link interface. |
| Last retransmission scan time is | Last retransmission scan time on the sham-link interface. |
| maximum is 0 msec | Maximum retransmission scan time on the sham-link interface. |
| Keychain-based authentication enabled | Keychain-based authentication is enabled. |
| Key id used is | Key ID used. |

show ospf summary-prefix

To display Open Shortest Path First (OSPF) aggregated summary address information, use the **show ospf summary-prefix** command in EXEC mode.

```
show ospf [ process-name ] [vrf {vrf-name| all}] summary-prefix
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |

Command Default

All summary prefixes

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf summary-prefix** command:

```
RP/0/RSP0/CPU0:router# show ospf summary-prefix
OSPF Process 1, summary-prefix
10.1.0.0/255.255.0.0 Metric 20, Type 2, Tag 0
```

This table describes the significant fields shown in the display.

Table 115: show ospf summary-prefix Field Descriptions

| Field | Description |
|----------------------|---|
| 10.1.0.0/255.255.0.0 | Summary address designated for a range of addresses. The IP subnet mask used for the summary route. |
| Metric | Metric used to advertise the summary routes. |
| Type | External link-state advertisements (LSA) metric type. |
| Tag | Tag value that can be used as a “match” value for controlling redistribution through route maps. |

Related Commands

| Command | Description |
|--|--|
| router ospf , on page 1078 | Configures an OSPF routing process. |
| summary-prefix (OSPF) , on page 1155 | Creates aggregate addresses for routes being redistributed from another routing protocol into the OSPF protocol. |

show ospf virtual-links

To display parameters and the current state of Open Shortest Path First (OSPF) virtual links, use the **show ospf virtual-links** command in EXEC mode.

```
show ospf [ process-name ] [vrf {vrf-name| all}] virtual-links
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | (Optional) Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | (Optional) Name of the OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings "default" and "all" are reserved VRF names. |
| all | (Optional) Specifies all OSPF VRF instances. |

Command Default

All virtual links

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospf virtual-links** command to display useful information for debugging OSPF routing operations.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospf virtual-links** command:

```
RP/0/RSP0/CPU0:router# show ospf virtual-links

Virtual Link to router 172.31.101.2 is up
Transit area 0.0.0.1, via interface GigabitEthernet 0/3/0/0, Cost of using 10
Transmit Delay is 1 sec, State POINT_TO_POINT
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

This table describes the significant fields shown in the display.

Table 116: show ospf virtual-links Field Descriptions

| Field | Description |
|---|--|
| Virtual Link to router 172.31.101.2 is up | OSPF neighbor and whether the link to that neighbor is up or down. |
| Transit area 0.0.0.1 | Transit area through which the virtual link is formed. |
| via interface GigabitEthernet 0/3/0/0 | Interface through which the virtual link is formed. |
| Cost of using using 10 | Cost of reaching the OSPF neighbor through the virtual link. |
| Transmit Delay is 1 sec | Transmit delay (in seconds) on the virtual link. |
| State POINT_TO_POINT | State of the OSPF neighbor. |
| Timer intervals | Various timer intervals (in seconds) configured for the link. |
| Hello due in 0:00:08 | When the next hello message is expected from the neighbor (in hh:mm:ss). |
| Adjacency State FULL | Adjacency state between the neighbors. |

Related Commands

| Command | Description |
|--|-------------------------------------|
| router ospf , on page 1078 | Configures an OSPF routing process. |

show protocols (OSPF)

To display information about the OSPFv2 processes running on the router, use the **show protocols** command in EXEC mode.

```
show protocols [afi-all| ipv4| ipv6] [all| protocol]
```

Syntax Description

| | |
|-----------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies an IPv4 address family. |
| ipv6 | (Optional) Specifies an IPv6 address family. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip For the IPv6 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospfv3 |

Command Default No default behavior or value

Command Modes EXEC

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |
| rib | read |

Examples

The following is an OSPF configuration and the resulting **show protocols ospf** display:

```
RP/0/RSP0/CPU0:router#show running router ospf 1
```

```
router ospf 1
router-id Loopback0
nsf
redistribute connected
redistribute isis 3
area 0
mpls traffic-eng
interface Loopback0
!
interface Loopback1
!
interface Loopback2
!
interface GigabitEthernet 0/3/0/0
!
interface GigabitEthernet 0/3/0/1
!
interface GigabitEthernet 0/3/0/2
!
interface GigabitEthernet 0/3/0/3
!
!
mpls traffic-eng router-id Loopback0
!
```

```
RP/0/RSP0/CPU0:router# show protocols ospf
Routing Protocol OSPF 1
Router Id: 55.55.55.55
Distance: 110
Non-Stop Forwarding: Enabled
Redistribution:
  connected
  isis 3
Area 0
MPLS/TE enabled
```

```

GigabitEthernet 0/3/0/3
GigabitEthernet 0/3/0/2
GigabitEthernet 0/3/0/1
GigabitEthernet 0/3/0/0
Loopback2
Loopback0

```

This table describes the significant fields shown in the display.

Table 117: show protocols ospf Field Descriptions

| Field | Description |
|---------------------|--|
| Router Id | ID of the router for this configuration. |
| Distance | Administrative distance of OSPF routes relative to routes from other protocols. |
| Non-Stop Forwarding | Status of nonstop forwarding. |
| Redistribution | Lists the protocols that are being redistributed. |
| Area | Information about the current area including list of interfaces and the status of Multiprotocol Label Switching traffic engineering (MPLS TE). |

snmp context (OSPF)

To specify an SNMP context for an OSPF instance, use the **snmp context** command in router configuration mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*

no snmp context *context_name*

Syntax Description

| | |
|---------------------|---|
| <i>context_name</i> | Specifies name of the SNMP context for OSPF instance. |
|---------------------|---|

Command Default

SNMP context is not specified.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **snmp-server** commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for information on using the **snmp-server** commands.



Note

To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPF protocol.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure an SNMP context *foo* for OSPF instance *100*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RSP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
1620
RP/0/RSP0/CPU0:router(config)#snmp-server community public RW
RP/0/RSP0/CPU0:router(config)#snmp-server contact foo
RP/0/RSP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPF instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo

router ospf 100
  router-id 2.2.2.2
  bfd fast-detect
  nsf cisco
  snmp context foo
  area 0
    interface Loopback1
    !
  !
  area 1
    interface GigabitEthernet0/2/0/1
      demand-circuit enable
    !
    interface POS0/3/0/0
    !
    interface POS0/3/0/1
    !
  !
  !
```

Related Commands

| Command | Description |
|---|---|
| snmp trap (OSPF) , on page 1149 | Enables SNMP trap for an OSPF instance |
| snmp-server host | Specifies the recipient of an SNMP notification operation. |
| snmp-server community | Configures the community access string to permit access to the Simple Network Management Protocol (SNMP). |
| snmp-server contact | Sets the Simple Network Management Protocol (SNMP) system contact. |

| Command | Description |
|----------------------------------|--|
| snmp-server community-map | Associates a Simple Network Management Protocol (SNMP) community with an SNMP context. |

snmp trap (OSPF)

To enable SNMP trap for an OSPF instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPF instance, use the **no** form of this command.

snmp trap

no snmp trap

Syntax Description This command has no keywords or arguments.

Command Default Disabled.

Command Modes VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.1.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | ospf | read, write |

Examples This example shows how to enable SNMP trap for OSPF instance *100* under VRF *vrf-1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#vrf vrf-1
RP/0/RSP0/CPU0:router(config-ospf-vrf)#snmp trap
```

| Related Commands | Command | Description |
|------------------|--|--|
| | snmp context (OSPF) , on page 1146 | Specifies SNMP context for an OSPF instance. |

snmp trap rate-limit (OSPF)

To control the number of traps that OSPF sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in router configuration mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*

no snmp trap rate-limit *window-size max-num-traps*

Syntax Description

| | |
|----------------------|--|
| <i>window-size</i> | Specifies the trap rate limit sliding window size. |
| <i>max-num-traps</i> | Specifies the maximum number of traps sent in window time. |

Command Default

None

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.9.0 | This command was introduced. This command replaces the snmp-server trap ospf rate-limit command. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|------------|
| ospf | read,write |

Examples

The following example shows how to set the trap rate limit sliding window size to 30 and the maximum number of traps sent to 100:

```
RP/0/RSP0/CPU0:router (config) #router ospf 100
RP/0/RSP0/CPU0:router (config-ospf) #snmp trap rate-limit 30 100
```

spf prefix-priority (OSPFv2)

To prioritize OSPFv2 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in router configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name*

no spf prefix-priority route-policy *policy-name*

Syntax Description

route-policy *policy-name* Specifies the route policy to apply to OSPFv2 prefix prioritization.

Note If SPF prefix prioritization is configured, /32 prefixes are no longer preferred by default. To retain the /32 prefixes in higher-priority queues, define the route-policy accordingly.

Command Default

SPF prefix prioritization is disabled.

Command Modes

OSPF router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

SPF prefix prioritization is disabled, by default. In disabled mode, the /32 prefixes are installed into the global RIB before other prefixes.

If SPF prefix prioritization is enabled, routes are matched against the route-policy criteria and are assigned to the appropriate priority queue based on the spf-priority set. Unmatched prefixes, including the /32 prefixes, are placed in the low-priority queue.

If all /32 prefixes are desired in the high-priority queue or medium-priority queue, configure the following single route map:

```
prefix-set ospf-medium-prefixes
  0.0.0.0/0 ge 32
end-set
```

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure OSPFv2 SPF prefix prioritization:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# prefix-set ospf-critical-prefixes
RP/0/RSP0/CPU0:router(config-pfx)# 66.0.0.0/16
RP/0/RSP0/CPU0:router(config-pfx)# end-set
RP/0/RSP0/CPU0:router(config)# route-policy ospf-spf-priority
RP/0/RSP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# router-id 66.0.0.1
RP/0/RSP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

Related Commands

| Command | Description |
|---------------------------|--|
| prefix-set | Enters prefix set configuration mode and defines a prefix set. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

stub (OSPF)

To define an area as a stub area, use the **stub** command in area configuration mode. To disable this function, use the **no** form of this command.

stub [no-summary]

no stub

Syntax Description

| | |
|-------------------|--|
| no-summary | (Optional) Prevents an Area Border Router (ABR) from sending summary link advertisements into the stub area. |
|-------------------|--|

Command Default

No stub area is defined.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **stub** command on all routers in the stub area.

Use the **default-cost** command on the ABR of a stub area to specify the cost of the default route advertised into the stub area by the ABR.

To further reduce the number of link-state advertisements (LSAs) sent into a stub area, you can configure the **no-summary** keyword on the ABR to prevent it from sending summary LSAs (LSA Type 3) into the stub area.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to assign a default cost of 20 to stub network 10.0.0.0:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
```

```
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0  
RP/0/RSP0/CPU0:router(config-ospf-ar)# stub  
RP/0/RSP0/CPU0:router(config-ospf-ar)# default-cost 20  
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/3
```

Related Commands

| Command | Description |
|--|---|
| authentication (OSPF), on page 936 | Enables authentication for an OSPF area. |
| default-cost (OSPF), on page 962 | Specifies a cost for the default summary route sent into a stub area. |

summary-prefix (OSPF)

To create aggregate addresses for routes being redistributed from another routing protocol into the Open Shortest Path First (OSPF) protocol, use the **summary-prefix** command in the appropriate mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *address mask* [**not-advertise** | **tag tag**]

no summary-prefix *address mask*

Syntax Description

| | |
|----------------------|---|
| <i>address</i> | Summary address designated for a range of addresses. |
| <i>mask</i> | IP subnet mask used for the summary route. |
| not-advertise | (Optional) Suppresses summary routes that match the address and mask pair from being advertised. |
| tag tag | (Optional) Tag value that can be used as a “match” value for controlling redistribution through route policies. |

Command Default

When this command is not used, specific addresses are created for each route from another route source being distributed into the OSPF protocol.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **summary-prefix** command to cause an OSPF Autonomous System Boundary Router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. This command summarizes only routes from other routing protocols that are being redistributed into OSPF.

You can use this command multiple times to summarize multiple groups of addresses. The metric used to advertise the summary is the lowest metric of all the more specific routes. This command helps reduce the size of the routing table.

If you want to summarize routes between OSPF areas, use the **range** command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

In the following example, summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

Related Commands

| Command | Description |
|---|---|
| range (OSPF) , on page 1065 | Consolidates and summarizes routes at an area boundary. |

timers lsa group-pacing

To change the interval at which Open Shortest Path First (OSPF) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, use the **timers lsa group-pacing** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa group-pacing *seconds*

no timers lsa group-pacing

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 seconds to 1800 seconds. |
|----------------|---|

Command Default

seconds : 240 seconds

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF LSA group pacing is enabled by default. For typical customers, the default group pacing interval for refreshing, checksumming, and aging is appropriate and you need not configure this feature.

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the OSPF pacing between LSA groups to 60 seconds:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf 1  
RP/0/RSP0/CPU0:router(config-ospf)# timers lsa group-pacing 60
```

timers lsa min-arrival

To limit the frequency that new instances of any particular Open Shortest Path First (OSPF) link-state advertisements (LSAs) can be accepted during flooding, use the **timers lsa min-arrival** command in the appropriate mode. To restore the default value, use the **no** form of this command.

timers lsa min-arrival *milliseconds*

no timers lsa min-arrival

Syntax Description

| | |
|---------------------|--|
| <i>milliseconds</i> | Minimum interval (in milliseconds) between accepting same LSA. Range is 0 to 600000 milliseconds. |
|---------------------|--|

Command Default

milliseconds : 100 milliseconds

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the minimum interval between accepting the same LSA to 2 seconds:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers lsa min-arrival 2
```

timers throttle lsa all (OSPF)

To modify the Open Shortest Path First (OSPF) link-state advertisement (LSA) throttling, use the **timers throttle lsa all** command in the appropriate mode. To revert LSA throttling to default settings, use the **no** form of this command

timers throttle lsa all *start-interval hold-interval max-interval*

no timers throttle lsa all

Syntax Description

| | |
|-----------------------|--|
| <i>start-interval</i> | Delay to generate first occurrence of LSA in milliseconds. Range is 0 to 600000 milliseconds. |
| <i>hold-interval</i> | Minimum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds. |
| <i>max-interval</i> | Maximum delay between originating the same LSA in milliseconds. Range is 1 to 600000 milliseconds. |

Command Default

start-interval : 50 milliseconds

hold-interval : 200 milliseconds

max-interval : 5000 milliseconds

Command Modes

Router configuration

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The *lsa-start* time is the delay before flooding the first instance of an LSA. The *lsa-hold* interval is the minimum time to elapse before flooding an updated instance of an LSA. The *lsa-max-wait* time is the maximum time that can elapse before flooding an updated instance of an LSA.

For quick convergence, use smaller times for the *lsa-start* time and *lsa-hold* interval. However, in relatively large networks, this may result in a large number of LSAs being flooded in a relatively short time. A balance with the *lsa-start* time and *lsa-hold* interval can be iteratively arrived at for the size of your network. The *lsa-max-wait* time can be used to ensure that OSPF reconverges within a reasonable amount of time.



Note LSA throttling is always enabled. You can change the timer values with the **timers throttle lsa all** command or specify the **no** keyword to revert back to the default settings.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 500, 1000, and 90,000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers throttle lsa all 500 1000 90000
```

The following example is output from the show ospf command that displays the modified LSA throttle settings:

```
RP/0/RSP0/CPU0:router# show ospf

Routing Process "ospf 1" with ID 1.1.1.1
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an area border router
  Initial SPF schedule delay 5000 msec
  Minimum hold time between two consecutive SPFs 10000 msec
  Maximum wait time between two consecutive SPFs 10000 msec
  Initial LSA throttle delay 500 msec
  Minimum hold time for LSA throttle 1000 msec
  Maximum wait time for LSA throttle 90000 msec
  Minimum LSA interval 1000 msec. Minimum LSA arrival 1 sec
  Maximum number of configured interfaces 255
  Number of external LSA 0. Checksum Sum 00000000
  Number of opaque AS LSA 0. Checksum Sum 00000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  External flood list length 0
  Non-Stop Forwarding enabled
    Area BACKBONE(0) (Inactive)
      Number of interfaces in this area is 2
      SPF algorithm executed 8 times
      Number of LSA 2. Checksum Sum 0x01ba83
      Number of opaque link LSA 0. Checksum Sum 00000000
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 0
      Flood list length 0
    Area 1
      Number of interfaces in this area is 1
      SPF algorithm executed 9 times
      Number of LSA 2. Checksum Sum 0x0153ea
      Number of opaque link LSA 0. Checksum Sum 00000000
      Number of DCbitless LSA 0
      Number of indication LSA 0
      Number of DoNotAge LSA 0
      Flood list length 0
```

Related Commands

| Command | Description |
|--|--|
| show ospf , on page 1084 | Displays generic information about OSPF routing processes. |

timers throttle spf (OSPF)

To modify the Open Shortest Path First (OSPF) shortest path first (SPF) throttling, use the **timers throttle spf** command in the appropriate mode. To revert SPF throttling to default settings, use the **no** form of this command.

timers throttle spf *spf-start spf-hold spf-max-wait*

no timers throttle spf

Syntax Description

| | |
|---------------------|--|
| <i>spf-start</i> | Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds. |
| <i>spf-hold</i> | Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |
| <i>spf-max-wait</i> | Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |

Command Default

spf-start:50 milliseconds
 spf-hold: 200 milliseconds
 spf-max-wait: 5000 milliseconds

Command Modes

Router configuration
 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The *spf-start* time is the delay before running SPF for the first time. The *spf-hold* interval is the minimum time to elapse between subsequent SPF runs. The *spf-max-wait* time is the maximum time that can elapse before running SPF again.



Tip

Setting a low *spf-start* time and *spf-hold* time causes routing to switch to the alternate path more quickly if there is a failure; however, it consumes more CPU processing time.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 5, 1000, and 90000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# timers throttle spf 5 1000 90000
```

transmit-delay (OSPF)

To set the estimated time required to send a link-state update packet on the interface, use the **transmit-delay** command in the appropriate mode. To return to the default value, use the **no** form of this command.

transmit-delay *seconds*

no transmit-delay *seconds*

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds. |
|----------------|--|

Command Default

seconds: 1 second

Command Modes

Router configuration
 Area configuration
 Interface configuration
 Virtual-link configuration
 VRF configuration
 Multi-area configuration
 Sham-link configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. This setting has significance only on very low-speed networks not supported in Cisco IOS XR software or on networks such as satellite circuits that incur a very long (greater than one second) delay time.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to configure a transmit delay for interface GigabitEthernet 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# area 0
RP/0/RSP0/CPU0:router(config-ospf-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospf-ar-if)# transmit-delay 3
```

Related Commands

| Command | Description |
|---|--|
| show ospf, on page 1084 | Displays general information about OSPF routing processes. |

virtual-link (OSPF)

To define an Open Shortest Path First (OSPF) virtual link, use the **virtual-link** command in area configuration mode. To remove a virtual link, use the **no** form of this command.

virtual-link *router-id*

no virtual-link *router-id*

Syntax Description

| | |
|------------------|--|
| <i>router-id</i> | Router ID associated with the virtual link neighbor. The router ID appears in the show ospf command display. The router ID can be any 32-bit router ID value specified in four-part, dotted-decimal notation. |
|------------------|--|

Command Default

No virtual links are defined.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All areas in an OSPF autonomous system must be physically connected to the backbone area (area 0). In some cases in which this physical connection is not possible, you can use a virtual link to connect to the backbone through a nonbackbone area. You can also use virtual links to connect two parts of a partitioned backbone through a nonbackbone area. The area through which you configure the virtual link, known as a transit area, must have full routing information. The transit area cannot be a stub or not-so-stubby area.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0
```

```
RP/0/RSP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RSP0/CPU0:router(config-ospf-ar-vl)#
```

The following example shows how to establish a virtual link with clear text authentication called mykey:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 201
RP/0/RSP0/CPU0:router(config-ospf)# area 10.0.0.0
RP/0/RSP0/CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
RP/0/RSP0/CPU0:router(config-ospf-ar-vl)# authentication-key 0 mykey
```

Related Commands

| Command | Description |
|---|---|
| authentication (OSPF), on page 936 | Enables authentication for an OSPF area. |
| show ospf virtual-links, on page 1141 | Displays parameters and the current state of OSPF virtual links |

vrf (OSPF)

To configure an Open Shortest Path First (OSPF) VPN routing and forwarding (VRF) instance, use the **vrf** command in router configuration mode. To terminate an OSPF VRF, use the **no** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

| | |
|-----------------|--|
| <i>vrf-name</i> | Identifier of an OSPF VRF. The <i>vrf-name</i> argument can be specified as an arbitrary string. The strings “default” and “all” are reserved VRF names. |
|-----------------|--|

Command Default

No OSPF VRF is defined.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to explicitly configure a VRF. Commands configured under the VRF configuration mode (such as the **interface** [OSPF] and **authentication** commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-id* argument format must be the same as the format used when creating the area.



Note

To remove the specified VRF from the router configuration, use the **no vrf** *vrf-id* command. The **no vrf** *vrf-id* command removes the VRF and all VRF options, such as **authentication**, **default-cost**, **nssa**, **range**, **stub**, **virtual-link**, and **interface**.

To avoid possibly having the router ID change under a VRF, explicitly configure the router ID using the **router-id** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure VRF vrf1 and GigabitEthernet interface 0/2/0/0. GigabitEthernet interface 0/2/0/0 is bound to VRF vrf1 automatically.

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# vrf vrf1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# interface GigabitEthernet 0/2/0/0
```

Related Commands

| Command | Description |
|---|---|
| router-id (OSPF) , on page 1076 | Configures a router ID for an OSPF process. |



OSPFv3 Commands

This module describes the commands used to configure and monitor the IP Version 6 (IPv6) Open Shortest Path First Version 3 (OSPFv3) routing protocol.

For detailed information about OSPFv3 concepts, configuration tasks, and examples, see the *Implementing OSPF on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family \(OSPFv3\)](#), page 1174
- [area \(OSPFv3\)](#), page 1175
- [authentication \(OSPFv3\)](#), page 1177
- [auto-cost \(OSPFv3\)](#), page 1179
- [capability vrf-lite \(OSPFv3\)](#), page 1181
- [clear ospfv3 process](#), page 1183
- [clear ospfv3 redistribution](#), page 1185
- [clear ospfv3 routes](#), page 1187
- [clear ospfv3 statistics](#), page 1189
- [cost \(OSPFv3\)](#), page 1191
- [database-filter all out \(OSPFv3\)](#), page 1193
- [dead-interval \(OSPFv3\)](#), page 1195
- [default-cost \(OSPFv3\)](#), page 1197
- [default-information originate \(OSPFv3\)](#), page 1199
- [default-metric \(OSPFv3\)](#), page 1201
- [demand-circuit \(OSPFv3\)](#), page 1203
- [distance ospfv3](#), page 1205
- [distribute-list prefix-list in](#), page 1207
- [distribute-list prefix-list out](#), page 1209
- [domain-id \(OSPFv3\)](#), page 1212

- encryption, page 1214
- flood-reduction (OSPFv3), page 1216
- graceful-restart (OSPFv3), page 1218
- hello-interval (OSPFv3), page 1220
- instance (OSPFv3), page 1222
- interface (OSPFv3), page 1224
- log adjacency changes (OSPFv3), page 1226
- maximum interfaces (OSPFv3), page 1228
- maximum paths (OSPFv3), page 1230
- maximum redistributed-prefixes (OSPFv3), page 1232
- mtu-ignore (OSPFv3), page 1234
- neighbor (OSPFv3), page 1236
- network (OSPFv3), page 1238
- nssa (OSPFv3), page 1240
- nsr (OSPFv3), page 1242
- ospfv3 name-lookup, page 1244
- packet-size (OSPFv3), page 1245
- passive (OSPFv3), page 1247
- priority (OSPFv3), page 1249
- protocol shutdown (OSPFv3), page 1251
- range (OSPFv3), page 1252
- redistribute (OSPFv3), page 1254
- retransmit-interval (OSPFv3), page 1259
- router-id (OSPFv3), page 1261
- router ospfv3, page 1263
- sham-link (OSPFv3), page 1265
- show ospfv3, page 1267
- show ospfv3 border-routers, page 1272
- show ospfv3 database, page 1274
- show ospfv3 flood-list, page 1287
- show ospfv3 interface, page 1290
- show ospfv3 message-queue, page 1293
- show ospfv3 neighbor, page 1295

- [show ospfv3 request-list](#), page 1302
- [show ospfv3 retransmission-list](#), page 1305
- [show ospfv3 routes](#), page 1308
- [show ospfv3 statistics rib-thread](#), page 1311
- [show ospfv3 summary-prefix](#), page 1313
- [show ospfv3 virtual-links](#), page 1315
- [show protocols \(OSPFv3\)](#), page 1317
- [snmp context \(OSPFv3\)](#), page 1319
- [snmp trap \(OSPFv3\)](#), page 1322
- [snmp trap rate-limit \(OSPFv3\)](#), page 1323
- [spf prefix-priority \(OSPFv3\)](#), page 1325
- [stub \(OSPFv3\)](#), page 1327
- [stub-router](#), page 1329
- [summary-prefix \(OSPFv3\)](#), page 1331
- [timers lsa arrival](#), page 1333
- [timers pacing flood](#), page 1335
- [timers pacing lsa-group](#), page 1337
- [timers pacing retransmission](#), page 1339
- [timers throttle lsa all \(OSPFv3\)](#), page 1341
- [timers throttle spf \(OSPFv3\)](#), page 1343
- [trace \(OSPFv3\)](#), page 1345
- [transmit-delay \(OSPFv3\)](#), page 1347
- [virtual-link \(OSPFv3\)](#), page 1349
- [vrf \(OSPFv3\)](#), page 1351

address-family (OSPFv3)

To enter address family configuration mode for Open Shortest Path First Version 3 (OSPFv3), use the **address-family** command in the router ospfv3 configuration mode. To disable address family configuration mode, use the **no** form of this command.

address-family ipv6 [unicast]

no address-family ipv6 [unicast]

Syntax Description

| | |
|----------------|---|
| ipv6 | Specifies IP Version 6 (IPv6) address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |

Command Default

An address family is not specified.

Command Modes

Router ospfv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the OSPFv3 router process with IPv6 unicast address prefixes:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# address-family ipv6 unicast
```

area (OSPFv3)

To configure an Open Shortest Path First Version 3 (OSPFv3) area, use the **area** command in an appropriate configuration mode. To remove an OSPFv3 area, use the **no** form of this command.

area *area-id*

no area *area-id*

Syntax Description

| | |
|----------------|---|
| <i>area-id</i> | Identifier of an OSPFv3 area. The <i>area-id</i> argument can be specified as either a decimal value or as an IPv4 address. |
|----------------|---|

Command Default

No OSPFv3 areas are defined.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

An area must be explicitly configured with the **area** command.

Use the **area** command to place the router in area configuration mode (prompt: config-router-ar), from which you can configure area-specific settings. Commands configured under this mode (such as the **interface** command) are automatically bound to that area.



Note

To remove the specified OSPFv3 area from the router ospfv3 configuration, use the **no area** *area-id* command. The **no area** *area-id* command removes the OSPFv3 area including all OSPFv3 area options, and all the OSPFv3 interfaces and interface options that are configured under the area.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ospf | read, write |

Examples

The following example shows how to configure area 0 for OSPFv3 process 1. The GigabitEthernet 0/1/0/1 interface also is configured:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
```

authentication (OSPFv3)

To enable plain text, Message Digest 5 (MD5) authentication, or null authentication for an Open Shortest Path First Version 3 (OSPFv3) interface, use the **authentication** command in an appropriate configuration mode. To remove such authentication, use the **no** form of this command.

```
authentication {ipsec spi spi-value {md5| sha1} [clear| password] password| disable}
```

```
no authentication
```

Syntax Description

| | |
|-----------------------------|--|
| ipsec | Specifies IP Security (IPSec). |
| spi <i>spi-value</i> | Specifies a security policy index (SPI) value. Range is 256 to 4294967295. |
| md5 | Enables Message Digest 5 (MD5) authentication. |
| sha1 | Enables SHA1 authentication. |
| clear | (Optional) Specifies that the key be unencrypted. |
| password | (Optional) Specifies that the key be encrypted using a two-way algorithm. |
| <i>password</i> | Any contiguous string that can be entered from the keyboard. |
| disable | Disables authentication for OSPFv3 packets. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the authentication parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the authentication parameter specified for the process.

If this command is not specified at any level, then the interface does not use authentication.

Command Modes

Interface configuration
 Area configuration
 Router configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **authentication** command to specify an authentication type for the interface, which overrides the authentication specified for the area to which this interface belongs. If this command is not included in the configuration file, the authentication configured in the area to which the interface belongs is assumed (as specified by the area **authentication** command).

The authentication type and password must be the same for all OSPFv3 interfaces that are to communicate with each other through OSPFv3.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable MD5 authentication:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.1.1.1
RP/0/RSP0/CPU0:router(config-ospfv3)# authentication ipsec spi 500 md5
1234567890abcdef1234567890abcdef
```

auto-cost (OSPFv3)

To control how the Open Shortest Path First Version 3 (OSPFv3) protocol calculates default metrics for an interface, use the **auto-cost** command in an appropriate configuration mode. To set link cost based only on the interface type, use the **disable** form of this command. To re-enable OSPFv3 metric calculation for an interface according to the bandwidth of the interface, use the **no** form of this command.

auto-cost [**reference-bandwidth** *mbps*| **disable**]

no auto-cost [**reference-bandwidth** *mbps*| **disable**]

Syntax Description

| | |
|--|--|
| reference-bandwidth <i>mbps</i> | (Optional) Sets the rate in Mbps (bandwidth). Range is 1 to 4294967. |
| disable | (Optional) Sets the link cost based only on the interface type. |

Command Default

mbps : 100 Mbps

Command Modes

Router ospfv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default OSPFv3 calculates the OSPFv3 metric for an interface according to the bandwidth of the interface.

The **no auto-cost disable** form of this command reenables OSPFv3 metric calculation for an interface according to the bandwidth of the interface.

To set link cost based only on the interface type, use the **disable** keyword.

If you have multiple links with high bandwidth, you might want to use a larger number to differentiate the cost on those links.

Recommended usage of cost configuration for all OSPFv3 configured interfaces is to be consistent: Either explicitly configure link costs (by using the **cost** command) or choose an appropriate default (by using the **auto-cost** command).

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the reference value for the auto cost to 64:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# auto-cost reference-bandwidth 64
```

Related Commands

| Command | Description |
|---|---|
| cost (OSPFv3), on page 1191 | Explicitly specifies the cost of the interface (network) for OSPF path calculation. |

capability vrf-lite (OSPFv3)

To ignore DN bit in LSAs received from peers in the given VRF and to disable automatic ABR status in that VRF, use the **capability vrf-lite** command in OSPFv3 VRF configuration mode. To disable ignoring the DN bit in LSAs and to re-enable automatic ABR status in the VRF, use the **no** form of this command.

capability vrf-lite

no capability vrf-lite

This command has no keywords or arguments.

Command Default

Disabled.

Command Modes

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **capability vrf-lite** command when routers (sometimes called multi-vrf CE routers) are directly connected through interfaces associated with a VRF, but without being connected to other PEs through the MPLS/VPN BGP Backbone.

When OSPFv3 is enabled in a VRF, the router is always an ABR. With the **capability vrf-lite** command enabled, the router becomes an ABR only if it is connected to area 0 (backbone area), and there are other (non-backbone) areas enabled on this router in the given VRF.



Note

Routes may be re-introduced to the VPN backbone when this command is used.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to enable VRF-lite capability for OSPFv3 instance *1* under VRF *vrf1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrf1
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#capability vrf-lite
```

Related Commands

| Command | Description |
|--|-------------------------------------|
| vrf (OSPFv3), on page 1351 | Enters OSPFv3 configuration submenu |

clear ospfv3 process

To reset an Open Shortest Path First Version 3 (OSPFv3) router process without removing and reconfiguring it, use the **clear ospfv3 process** command in EXEC configuration mode.

```
clear ospfv3 [process-name] [vrf vrf-name] process
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf-name keyword and argument were added to support OSPFv3 VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the OSPFv3 router process is reset, OSPFv3 releases all resources allocated, cleans up the internal database, uninstalls routes, and resets all OSPFv3 adjacencies.



Note

The **clear ospfv3 process** command might clear the router ID configuration if the OSPF router ID is not explicitly configured through the [router-id \(OSPFv3\)](#), [on page 1261](#) command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to reset all OSPFv3 processes:

```
RP/0/RSP0/CPU0:router# clear ospfv3 process
```

The following example shows how to reset the OSPFv3 process 1:

```
RP/0/RSP0/CPU0:router# clear ospfv3 1 process
```

Related Commands

| Command | Description |
|--|--|
| router-id (OSPFv3), on page 1261 | Configures a router ID for the OSPFv3 process. |

clear ospfv3 redistribution

To flush all the Type 5 and Type 7 link-state advertisements (LSAs) originated by an Open Shortest Path First Version 3 (OSPFv3) process, use the **clear ospfv3 redistribution** command in EXEC configuration mode.

clear ospfv3 [*process-name*] [**vrf** *vrf-name*] **redistribution**

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf <i>vrf-name</i> keyword and argument were added to support OSPFv3 VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 redistribution** command to cause the routing table to be read again. OSPFv3 regenerates and sends Type 5 and Type 7 link-state advertisements (LSAs) to its neighbors. If an unexpected route has appeared in the OSPFv3 redistribution, using this command corrects the issue.



Note

Use of this command can cause a significant number of LSAs to flood the network. We recommend that you use this command with caution.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to clear all OSPFv3 redistributed routes from other protocols:

```
RP/0/RSP0/CPU0:router# clear ospfv3 redistribution
```

clear ospfv3 routes

To clear the Open Shortest Path First Version 3 (OSPFv3) internal route table, use the **clear ospfv3 routes** command in EXEC configuration mode.

clear ospfv3 [*process-name*] [**vrf** *vrf-name*] **routes**

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only the specified routing process is affected. Otherwise, all OSPFv3 processes are reset. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf <i>vrf-name</i> keyword and argument were added to support OSPFv3 VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 routes** command to force the internal route table to be repopulated by causing recalculation of the shortest path first (SPF) routing table. When the OSPFv3 routing table is cleared, OSPFv3 routes in the global routing table are also recalculated.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to clear all OSPFv3 routes from the OSPFv3 routing table and recomputes valid routes:

```
RP/0/RSP0/CPU0:router# clear ospfv3 routes
```

clear ospfv3 statistics

To clear the Open Shortest Path First Version 3 (OSPFv3) statistical counters, use the **clear ospfv3 statistics** command in EXEC configuration mode.

```
clear ospfv3 [process-name] [vrf vrf-name] statistics [neighbor [type interface-path-id] [router-id]]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only the specified routing process is affected. |
| neighbor | (Optional) Clears counters for the specified neighbor only. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>router-id</i> | (Optional) Specified router ID. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address. This argument clears the counters of the specified neighbor only. |
| vrf | (Optional) Specifies VPN routing and forwarding (VRF). |
| <i>vrf-name</i> | Name of a VRF. |

Command Default

No default behavior or value

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf-name keyword and argument were added to support OSPFv3 VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear ospfv3 statistics** command to reset statistics so that subsequent changes are easily observed.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to clear the OSPFv3 statistical counters of all neighbors on GigabitEthernet interface 0/2/0/0:

```
RP/0/RSP0/CPU0:router# clear ospfv3 statistics neighbor GigabitEthernet 0/2/0/0
```

cost (OSPFv3)

To explicitly specify the cost of the interface (network) for OSPF path calculations, use the **cost** command in an appropriate configuration mode. To remove the cost, use the **no** form of this command.

cost *cost*

no cost

Syntax Description

| | |
|-------------|---|
| <i>cost</i> | Unsigned integer value expressed as the link-state metric. Range is 1 to 65535. |
|-------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the cost parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the cost parameter specified for the process.

If this command is not specified at any level, then the cost is based on the interface bandwidth, as specified by the **auto-cost** command.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The link-state metric is advertised as the link cost in the router link advertisement.

In general, the path cost is calculated using the following formula:

$$10^8 / \text{bandwidth}$$

If this value does not suit your network, you can use your own method of calculating path costs.

The value set by the **cost** command overrides the cost resulting from the **auto-cost** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the cost value to 65 for GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 65
```

Related Commands

| Command | Description |
|--|---|
| auto-cost (OSPFv3), on page 1179 | Controls how the OSPFv3 protocol calculates default metrics for an interface. |

database-filter all out (OSPFv3)

To filter outgoing link-state advertisements (LSAs) to an Open Shortest Path First Version 3 (OSPFv3) interface, use the **database-filter all out** command in an appropriate configuration mode. To restore the forwarding of LSAs to the interface, use the **no** form of this command.

database-filter all out

no database-filter all out

Syntax Description This command has no keywords or arguments.

Command Default If this command is not specified in interface configuration mode, then the interface adopts the database filter parameter specified for the area.
If this command is not specified in area configuration mode, then the interface adopts the database filter parameter specified for the process.
If this command is not specified in router ospfv3 configuration mode, then the database filter is disabled and all outgoing LSAs are flooded to the interface.

Command Modes Interface configuration
Area configuration
Router OSPFv3 configuration
OSPFv3 VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|--|
| | Release 3.7.2 | This command was introduced. |
| | Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **database-filter all out** command to perform the same function that the **neighbor** command (with the **database-filter** keyword) performs on a neighbor basis.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to prevent flooding of OSPFv3 LSAs to neighbors reachable through GigabitEthernet interface 0/2/0/3:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# database-filter all out
```

Related Commands

| Command | Description |
|---|---|
| neighbor (OSPFv3), on page 1236 | Configures OSPFv3 routers interconnecting to nonbroadcast networks. |

dead-interval (OSPFv3)

To set the interval after which a neighbor is declared dead when no hello packets are observed, use the **dead-interval** command in an appropriate configuration mode. To return to the default time, use the **no** form of this command.

dead-interval *seconds*

no dead-interval

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Unsigned integer that specifies the interval (in seconds). The value must be the same for all nodes on the same network link. Range is 1 to 65535. |
|----------------|--|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the dead interval parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the dead interval parameter specified for the process.

If this command is not specified in router ospfv3 configuration mode, then the dead interval is four times the interval set by the **hello-interval (OSPFv3)** command.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Two Open Shortest Path First Version 3 (OSPFv3) routers do not become adjacent if their dead interval values differ.

If the hello interval is configured, the dead interval value must be larger than the hello interval value. The dead interval value is usually configured four times larger than the hello interval value.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the OSPFv3 dead interval on GigabitEthernet interface 0/2/0/3 to 40 seconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# dead-interval 40
```

Related Commands

| Command | Description |
|---|---|
| hello-interval (OSPFv3), on page 1220 | Specifies the interval between hello packets that the Cisco IOS XR software sends on the interface. |

default-cost (OSPFv3)

To specify a cost for the default summary route sent into a stub area or not-so-stubby area (NSSA) for Open Shortest Path First Version 3 (OSPFv3) packets, use the **default-cost** command in area configuration mode. To remove the assigned default route cost, use the **no** form of this command.

default-cost *cost*

no default-cost

Syntax Description

| | |
|-------------|--|
| <i>cost</i> | Cost for the default summary route used for a stub or NSSA area. The acceptable value is a 24-bit number ranging from 1 to 16777214. |
|-------------|--|

Command Default

cost : 1

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-cost** command only on an Area Border Router (ABR) attached to a stub or an NSSA area.

In all routers and access servers attached to the stub area, the area should be configured as a stub area using the **stub (OSPFv3)** command in the area configuration submenu. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to assign a cost of 20 to the default route sent into area 10.15.0.0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 10.15.0.0
```

```
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# stub
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# default-cost 20
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
```

Related Commands

| Command | Description |
|---|---------------------------------|
| stub (OSPFv3), on page 1327 | Defines an area as a stub area. |

default-information originate (OSPFv3)

To generate a default external route into an Open Shortest Path First Version 3 (OSPFv3) routing domain, use the **default-information originate** command in router ospfv3 configuration mode. To disable this feature, use the **no** form of this command.

default-information originate [*route-policy policy-name*] [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**tag** *tag-value*]

no default-information originate [*route-policy policy-name*] [**always**] [**metric** *metric-value*] [**metric-type** *type-value*] [**tag** *tag-value*]

Syntax Description

| | |
|--|---|
| route-policy <i>policy-name</i> | (Optional) Specifies the route policy to apply to default information origination. |
| always | (Optional) Always advertises the default route regardless of whether the software has a default route. |
| metric <i>metric-value</i> | (Optional) Specifies a metric used for generating the default route. The default metric value is 1. The value used is specific to the protocol. |
| metric-type <i>type-value</i> | (Optional) Specifies an external link type associated with the default route advertised into the OSPFv3 routing domain. It can be one of the following values: 1—Type 1 external route 2—Type 2 external route |
| tag <i>tag-value</i> | (Optional) 32-bit dotted-decimal value attached to each external route. This is not used by the OSPFv3 protocol itself. It may be used to communicate information between autonomous system boundary routers (ASBRs). If a tag is not specified, then zero (0) is used. |

Command Default

A default external route into an OSPFv3 routing domain is not generated.

metric-value : 1

type-value : Type 2

tag-value: 0

Command Modes

Router ospfv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Whenever you use the **redistribute** or **default-information** command to redistribute routes into an OSPFv3 routing domain, the software automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPFv3 routing domain. The software still must have a default route for itself before it generates one, except when you have specified the **always** keyword.

The **default-information originate** route-policy attach point conditionally injects the default route 0.0.0.0/0 into the OSPF link-state database, and is done by evaluating the attached policy. If any routes specified in the policy exist in the global RIB, then the default route is inserted into the link-state database. If there is no match condition specified in the policy, the policy passes and the default route is generated into the link-state database.

For information about the default-information originate attach point, see the *OSPF v3 Policy Attach Points* section in the *Implementing Routing Policy* chapter in *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

When you use the **default-information originate** command for the OSPFv3 process, the default network must reside in the routing table.

For information about routing policies, see the *Routing Policy Commands* chapter in the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify a metric of 100 for the default route redistributed into the OSPFv3 routing domain and an external metric type of Type 1:

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)#default-information originate metric 100 metric-type 1
```

Related Commands

| Command | Description |
|---|---|
| redistribute (OSPFv3), on page 1254 | Redistributes routes from one routing domain into another routing domain. |

default-metric (OSPFv3)

To set default metric values for routes redistributed from another protocol into Open Shortest Path First Version 3 (OSPFv3), use the **default-metric** command in an appropriate configuration mode. To return to the default state, use the **no** form of this command.

default-metric *value*

no default-metric *value*

Syntax Description

| | |
|--------------|--|
| <i>value</i> | Default metric value appropriate for the specified routing protocol. |
|--------------|--|

Command Default

Built-in, automatic metric translations, as appropriate for each routing protocol

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submodule. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause the current routing protocol to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics. Whenever metrics do not convert, use a default metric to provide a reasonable substitute and enable the redistribution to proceed.

The default-metric value configured in OSPF configuration does not apply to connected routes that are redistributed to OSPF using the **redistribute connected** command. To set a non-default metric for connected routes, configure OSPF with the **redistribute connected metric** *metric-value* command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure a router with both the Intermediate System-to-Intermediate System (IS-IS) and the OSPFv3 routing protocols. The OSPFv3 routing protocol advertises IS-IS derived routes and assigns the routes a metric of 10:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# default-metric 10
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute isis IS-IS_isp
```

Related Commands

| Command | Description |
|---|---|
| redistribute (OSPFv3), on page 1254 | Redistributes routes from one routing domain into another routing domain. |

demand-circuit (OSPFv3)

To configure the Open Shortest Path First Version 3 (OSPFv3) router process to treat the interface as an OSPFv3 demand circuit, use the **demand-circuit** command in an appropriate configuration mode. To remove the demand circuit designation from the interface, use the **no** form of this command.

demand-circuit [disable]

no demand-circuit

Syntax Description

| | |
|----------------|---|
| disable | (Optional) Disables the demand circuit configuration that may have been specified at a higher level in the configuration. |
|----------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the demand circuit parameter specified for the area.

If this command is not specified in area configuration mode, then the interface adopts the demand circuit parameter specified for the process.

If this command is not specified at any level, then the interface is not a demand circuit.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

On point-to-point interfaces, only one end of the demand circuit must be configured with the **demand-circuit** command. Periodic hello messages are suppressed and periodic refreshes of link-state advertisements (LSAs) do not flood the demand circuit. This command allows the underlying data link layer to be closed when the

topology is stable. In point-to-multipoint topology, only the multipoint end must be configured with this command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure GigabitEthernet interface 0/3/0/1 as an on-demand circuit:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0  
RP/0/RSP0/CPU0:router(config-ospfv3)# interface GigabitEthernet 0/3/0/1  
RP/0/RSP0/CPU0:router(config-ospfv3-if)# demand-circuit
```

distance ospfv3

To define the Open Shortest Path First Version 3 (OSPFv3) route administrative distances based on route type, use the **distance ospfv3** command in router ospfv3 configuration mode. To restore the default value, use the **no** form of this command.

distance ospfv3 {**intra-area**|**inter-area**|**external**} *distance*

no distance ospfv3

Syntax Description

| | |
|---|---|
| intra-area inter-area external | Type of area. It can be one of the following values: intra-area —All routes within an area. inter-area —All routes from one area to another area. external —All routes from other routing domains, learned by redistribution. |
| <i>distance</i> | The route administrative distance. |

Command Default

distance : 110

Command Modes

Router ospfv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must specify one of the keywords.

Use the **distance ospfv3** command to perform the same function as the **distance** command used with an access list. However, the **distance ospfv3** command sets a distance for an entire group of routes, rather than a specific route that passes an access list.

A common reason to use the **distance ospfv3** command is when you have multiple OSPFv3 processes with mutual redistribution, and you want to prefer internal routes from one over external routes from the other.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the external distance to 200, making it less reliable:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute ospfv3 2  
RP/0/RSP0/CPU0:router(config-ospfv3)# distance ospfv3 external 200  
RP/0/RSP0/CPU0:router(config-ospfv3)# exit  
RP/0/RSP0/CPU0:router(config)# router ospfv3 2  
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# distance ospfv3 external 200
```

Related Commands

| Command | Description |
|---------------|-------------------------------------|
| distance ospf | Defines an administrative distance. |

distribute-list prefix-list in

To filter the routes that Open Shortest Path First Version 3 (OSPFv3) installs in the Routing Information Base (RIB), use the **distribute-list prefix-list in** command in an appropriate configuration mode. To remove the filter, use the **no** form of this command.

distribute-list prefix-list *prefix-list-name* **in**

no distribute-list prefix-list *prefix-list-name* **in**

Syntax Description

| | |
|-------------------------|--|
| <i>prefix-list-name</i> | IP Version 6 (IPv6) prefix list name. The list defines which IPv6 prefixes are installed in the RIB. |
|-------------------------|--|

Command Default

All routes learned by OSPFv3 are installed in the RIB.

Command Modes

Interface configuration
Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **distribute-list prefix-list** command to limit the routes that OSPFv3 installs in the RIB of your router. This command does not affect the information sent to other OSPFv3 routers or the routes that these routers compute and install.

**Note**

Because the other OSPFv3 routers are not aware of any omissions in the RIB, they may send traffic addressed to the missing prefixes. If no other provision has been made for these prefixes, the packets are dropped.

When this command is specified in router ospfv3 configuration mode, the filter applies to all routes computed by OSPFv3.

When this command is specified in interface configuration mode, the filter applies only to routes that forward outgoing traffic over that interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to prevent OSPFv3 from installing any routes that have 2001:e624 as the first 32 bits of the address. OSPFv3 is also prevented from installing routes to 2002::/16 that use GigabitEthernet interface 0/2/0/0 as the next-hop interface:

```
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list preflist1
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2001:e624::/32 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list preflist2
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2002::/16
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list preflist1 in
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# distribute-list prefix-list preflist2 in
```

distribute-list prefix-list out

To filter the routes redistributed into Open Shortest Path First Version 3 (OSPFv3) from other routing protocols, use the **distribute-list prefix-list out** command in an appropriate configuration mode. To remove the filter, use the **no** form of this command.

distribute-list prefix-list *prefix-list-name* **out** [*protocol* [*process-id*]]

no distribute-list prefix-list *prefix-list-name* **out** [*protocol* [*process-id*]]

Syntax Description

prefix-list-name IP Version 6 (IPv6) prefix list name. The list defines which IPv6 prefixes are installed in the RIB.

protocol (Optional) Source protocol from which routes are being redistributed. It can be one of the following keywords: **bgp**, **eigrp**, **isis**, **ospfv3**, **static**, and **connected**.

The **static** keyword is used to redistribute IPv6 static routes.

The **connected** keyword refers to routes that are established automatically because IPv6 is enabled on an interface. For routing protocols such as OSPFv3 and Intermediate System-to-Intermediate System (IS-IS), these routes are redistributed as external to the autonomous system.

process-id (Optional) For the **bgp** keyword, an autonomous system number has the following ranges:

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

For the **eigrp** keyword, an autonomous system number.

For the **isis** keyword, an optional argument that defines a meaningful name for a routing process. You can specify only one IS-IS process for each router. Creating a name for a routing process means that you use names when configuring routing.

For the **ospfv3** keyword, an appropriate OSPFv3 process name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.

Command Default

All routes from protocols specified in the [redistribute \(OSPFv3\)](#), on page 1254 command are redistributed into OSPFv3.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Routes may be redistributed into OSPFv3 from several other routing protocols or from other OSPFv3 processes. These routes are then communicated to other OSPFv3 routes through Type 5 (External) or Type 7 not-so-stubby area (NSSA) link-state advertisements (LSAs). Use the **distribute-list prefix-list out** command to control redistribution by matching redistributed routes against an IPv6 prefix list. Only routes permitted by the prefix list are redistributed into OSPFv3.

Each protocol being redistributed into OSPFv3 can have a separate prefix list. In addition, a prefix list can be defined that applies to all protocols.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to prevent OSPFv3 from redistributing routes that have 2001:e624 as the first 32 bits of the address. In addition, routes with a prefix beginning with 2064 are not redistributed from Border Gateway Protocol (BGP) autonomous system 1, and only those routes are redistributed from BGP autonomous system 5.

```
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p1
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2001:e624::/32 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p2
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# deny 2064::/16 le 128
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit ::/0 le 128
!
RP/0/RSP0/CPU0:router(config)# ipv6 prefix-list p3
RP/0/RSP0/CPU0:router(config-ipv6-pfx)# permit 2064::/16 le 128
!
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 5
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p1 out
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p2 out bgp 1
RP/0/RSP0/CPU0:router(config-ospfv3)# distribute-list prefix-list p3 out bgp 5
```

Related Commands

| Command | Description |
|---|--|
| redistribute (OSPFv3), on page 1254 | Redistributes routes from one routing domain into another routing domain for OSPFv3. |

domain-id (OSPFv3)

To specify the Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) domain ID, use the **domain-id** command in VRF configuration mode. To remove the OSPFv3 VRF domain ID, use the **no** form of this command.

domain-id [**secondary**] **type** [0005| 0105| 0205] **value** *domain-id_value*

no domain-id [**secondary**] **type** [0005| 0105| 0205] **value** *domain-id-value*

Syntax Description

| | |
|------------------------|--|
| secondary | (Optional) OSPFv3 secondary domain ID. |
| type | Primary OSPFv3 domain ID in hexadecimal format. <ul style="list-style-type: none"> • 0005 —Type 0x0005 • 0105 —Type 0x0105 • 0205—Type 0x0205 |
| value | OSPF domain ID value in hexadecimal format. |
| <i>domain-id-value</i> | OSPF domain ID extended community value as a 6 byte hexadecimal number. |

Command Default

No domain ID is specified.

Command Modes

VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If no value is specified for the domain ID, the default is Null (all zeros) primary domain ID. One or more secondary domain IDs can be specified.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to specify a domain ID with type *0105* and value *AABBCCDDEEFF*:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospf 1
RP/0/RSP0/CPU0:router(config-ospf)# vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospf-vrf)# domain-id type 0105 value AABBCCDDEEFF
```

Related Commands

| Command | Description |
|--|-------------------------------------|
| vrf (OSPFv3), on page 1351 | Enters OSPFv3 configuration submenu |

encryption

To encrypt and authenticate Open Shortest Path First Version 3 (OSPFv3) packets, use the **encryption** command in an appropriate configuration mode. To remove the encryption, use the **no** form of this command.

```
encryption {disable| ipsec spi spi-value esp {3des| aes [192| 256] | des| null [clear| password]
encrypt-password} [authentication {md5| sha1} [clear| password] auth-password}
```

no encryption

Syntax Description

| | |
|-------------------------|--|
| disable | Disables OSPFv3 packet encryption. |
| ipsec spi | Specifies IPsec ESP encryption and authentication with the Security Parameter Index (SPI) value. |
| <i>spi-value</i> | SPI value. Range is 256 to 4294967295. |
| esp | Specifies Encryption Service Payload (ESP) encryption parameters. |
| 3des | Specifies the triple DES algorithm. |
| aes | Specifies the Advanced Encryption Standard (AES) algorithm. |
| 192 | (Optional) Specifies the 192-bit AES algorithm |
| 256 | (Optional) Specifies the 256-bit AES algorithm |
| des | Specifies the Data Encryption Standard (DES) algorithm. |
| null | Specifies no AES algorithm. |
| md5 | Enables Message Digest 5 (MD5) authentication. |
| sha1 | Enables SHA1 authentication. |
| clear | Specifies that the key be unencrypted. |
| password | Specifies that the key be encrypted using a two-way algorithm. |
| <i>encrypt-password</i> | Any contiguous string that can be entered from the keyboard as the encryption password. |
| <i>auth-password</i> | Any contiguous string that can be entered from the keyboard as the authentication password. |

Command Default

No default behavior or values.

Command Modes

Interface configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|--|
| | Release 3.7.2 | This command was introduced. |
| | Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **encryption** command to encrypt and authenticate OSPFv3 packets.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to encrypt and authenticate OSPFv3 packets:

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)#encryption ipsec spi 256 esp 3des clear
```

flood-reduction (OSPFv3)

To suppress the unnecessary flooding of link-state advertisements (LSAs) in stable topologies, use the **flood-reduction** command in an appropriate configuration mode. To disable this feature, use the **no** form of this command.

flood-reduction [disable]

no flood-reduction

Syntax Description

| | |
|----------------|--|
| disable | (Optional) Turns off this functionality at a specific level. |
| Note | The disable keyword is not available in router ospfv3 configuration mode. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the flood reduction parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the flood reduction parameter specified for the process.

If this command is not specified at any level, then flood reduction is disabled.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All routers that support Open Shortest Path First Version 3 (OSPFv3) demand circuits are compatible with routers supporting flooding reduction.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to reduce the flooding of unnecessary LSAs for area 0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# flood-reduction
```

Related Commands

| Command | Description |
|---|--|
| show ospfv3 interface, on page 1290 | Displays OSPFv3-related interface information. |
| show ospfv3 neighbor, on page 1295 | Displays OSPFv3 neighbor information on an individual interface basis. |

graceful-restart (OSPFv3)

To enable graceful restart, use the **graceful-restart** command in an appropriate configuration mode. To disable this feature, use the **no** form of this command.

graceful-restart [**helper disable**| **interval** *interval*| **lifetime** *lifetime*]

no graceful-restart [**helper disable**| **interval** *interval*| **lifetime** *lifetime*]

Syntax Description

| | |
|---------------------------------|---|
| helper disable | (Optional) Disables the routers helper support level. |
| interval <i>interval</i> | (Optional) Specifies the minimum interval between graceful restarts. Range is 90 to 3600 seconds. |
| lifetime <i>lifetime</i> | (Optional) Specifies the maximum route lifetime following a restart. Range is 90 to 3600 seconds. |

Command Default

No default behavior or values.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to enable the Graceful Restart feature with a minimum interval between restarts of 300 seconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# graceful-restart interval 300
```

Related Commands

| Command | Description |
|---|--|
| show ospfv3 interface, on page 1290 | Displays OSPFv3-related interface information. |
| show ospfv3 neighbor, on page 1295 | Displays OSPFv3 neighbor information on an individual interface basis. |

hello-interval (OSPFv3)

To specify the interval between hello packets that Open Shortest Path First Version 3 (OSPFv3) sends on an interface, use the **hello-interval** command in an appropriate configuration mode. To return to the default time, use the **no** form of this command.

hello-interval *seconds*

no hello-interval

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Interval (in seconds). The value must be the same for all nodes on a specific network. |
|----------------|--|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the hello interval parameter specified by area.

If this command is not specified in area configuration mode, then the interface adopts the hello interval parameter specified for the process.

If this command is not specified at any level, then the hello interval is 10 seconds (broadcast) or 30 seconds (nonbroadcast).

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submodule. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The hello interval value is advertised in the hello packets. The shorter the hello interval, the faster topological changes are detected, but more routing traffic ensues. This value must be the same for all routers and access servers on a specific network.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the interval between hello packets to 15 seconds on GigabitEthernet interface 0/3/0/2:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# hello-interval 15
```

Related Commands

| Command | Description |
|--|---|
| dead-interval (OSPFv3), on page 1195 | Sets the interval after which a neighbor is declared dead when no hello packets are observed. |

instance (OSPFv3)

To set the 8-bit instance ID used in Open Shortest Path First Version 3 (OSPFv3) packets sent on an interface, use the **instance** command in an appropriate configuration mode. To remove the instance ID, use the **no** form of this command.

instance *instance-id*

no instance *instance-id*

Syntax Description

| | |
|--------------------|---|
| <i>instance-id</i> | Instance identifier sent in OSPFv3 packets. Range is 0 to 255. The same value must be used by all the communicating OSPFv3 routers on a link. |
|--------------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the instance parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the instance parameter specified for the process.

If this command is not specified at any level, then the instance is 0.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The OSPFv3 routing protocol allows multiple unrelated OSPFv3 processes to share a link by using an 8-bit "instance" value to demultiplex the protocol packets. Each OSPFv3 process sets its configured instance value in the OSPFv3 packets that it sends, and ignores received packets with instance values from other OSPFv3 processes.

**Note**

The *instance-id* argument should not be confused with the *process-name* argument that is specified by the **router ospfv3** command. The former is an 8-bit integer that is sent to other routers as part of the OSPFv3 protocol, and the latter is a 1- to 40-character ASCII string that is significant only within a given router. The instance ID value is also unrelated to the router ID that is specified by the **router-id** command, which is a 32-bit integer value that uniquely identifies a router within an OSPFv3 routing domain.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the instance value for GigabitEthernet interface 0/3/0/1 to 42:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# instance 42
```

Related Commands

| Command | Description |
|---|--|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |
| router-id (OSPFv3) , on page 1261 | Configures a router ID for the OSPFv3 routing process. |

interface (OSPFv3)

To define the interfaces on which Open Shortest Path First Version 3 (OSPFv3) runs, use the **interface** command in an appropriate configuration mode. To disable OSPFv3 routing for an interface, use the **no** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|---|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

An interface is not defined.

Command Modes

Area configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with an OSPFv3 area. The interface remains associated with the area even when the IPv6 address of the interface changes.

Similar to IPv4 address behavior for the **interface** command, all configured IPv6 addresses are advertised on an interface after the interface is associated to the OSPF routing process. The only difference is, IPv6 addresses can have multiple primary addresses.

This command places the router in interface configuration mode (prompt: config-router-ar-if), from which you can configure interface-specific settings. Commands configured under this mode (such as the **cost** command) are automatically bound to that interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to define two interfaces that belong to area 1. The cost value for packets on GigabitEthernet interface 0/3/0/1 is set at 40; the cost value for GigabitEthernet interface 0/3/0/2 is 65:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 40
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# cost 65
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
```

log adjacency changes (OSPFv3)

To change the default syslog messages for Open Shortest Path First Version 3 (OSPFv3) neighbor state changes, use the **log adjacency changes** command in an appropriate configuration mode. To suppress all adjacency change messages, use the **disable** keyword.

log adjacency changes [**detail**| **disable**]

Syntax Description

| | |
|----------------|---|
| detail | (Optional) Provides all (DOWN, INIT, 2WAY, EXSTART, EXCHANGE, LOADING, FULL) adjacency state changes. |
| disable | (Optional) Disables the neighbor state change messages. |

Command Default

Neighbor state change messages are enabled.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

By default, you are notified of OSPFv3 neighbor changes without explicitly configuring the **log adjacency changes** command. The syslog messages that are sent provide a high-level view of changes to the state of the OSPFv3 peer relationship.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to disable neighbor state change messages:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# log adjacency changes disable
```

The following example shows how to re-enable syslog messages for any OSPFv3 neighbor state changes:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# log adjacency changes
```

maximum interfaces (OSPFv3)

To control the maximum number of interfaces that can be configured under an Open Shortest Path First Version 3 (OSPFv3) process, use the **maximum interfaces** command in an appropriate configuration mode. To remove the **maximum interfaces** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum interfaces *number-interfaces*

no maximum interfaces

Syntax Description

| | |
|--------------------------|--|
| <i>number-interfaces</i> | Maximum number of interfaces that can be configured for this OSPFv3 process. Range is 1 to 4294967295. |
|--------------------------|--|

Command Default

If the command is not specified, the default is 1024.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.1 | The range for number of interfaces was changed to 1 to 4294967295 from 1 to 1024. The default number of interfaces was changed to 1024 from 255. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to allow a maximum of 1500 interfaces in an OSPFv3 process:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum interfaces 1500
```

maximum paths (OSPFv3)

To control the maximum number of parallel routes that the Open Shortest Path First Version 3 (OSPFv3) can support, use the **maximum paths** command in an appropriate configuration mode. To remove the **maximum paths** command from the configuration file and restore the system to its default condition with respect to the routing protocol, use the **no** form of this command.

maximum paths *maximum-routes-number*

no maximum paths

Syntax Description

| | |
|------------------------------|---|
| <i>maximum-routes-number</i> | Maximum number of parallel routes that OSPFv3 can install in a routing table. Range is 1 to 32. |
| Note | The maximum number of paths that can be configured is 32. |

Command Default

32 paths

Command Modes

Router OSPFv3 configuration
VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the maximum number of parallel routes is reduced, all existing paths are pruned and paths reinstalled at the new maximum number. During this route-reduction period, you may experience some packet loss for a few seconds. This may impact route traffic.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to allow a maximum of two paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum paths 2
```

maximum redistributed-prefixes (OSPFv3)

To limit the number of prefixes redistributed into Open Shortest Path First Version 3 (OSPFv3) or to generate a warning when the number of prefixes redistributed into OSPFv3 reaches a maximum, use the **maximum redistributed-prefixes** command in an appropriate configuration mode. To remove the values, use the **no** form of this command.

maximum redistributed-prefixes *limit* [*threshold*] [**warning-only**]

no maximum redistributed-prefixes

Syntax Description

| | |
|---------------------|--|
| <i>limit</i> | Maximum number of IP Version 6 (IPv6) prefixes that are allowed to be redistributed into OSPFv3, or, if the warning-only keyword is present, sets the number of prefixes allowed to be redistributed into OSPFv3 before the system logs a warning message. Range is 1 to 4294967295. Note If the warning-only keyword is also configured, this value does not limit redistribution; it is simply the number of redistributed prefixes that, when reached, causes a warning message to be logged. |
| <i>threshold</i> | (Optional) Percentage of the value set for the maximum number of redistributed prefixes that, when reached, causes a warning message to be logged. |
| warning-only | (Optional) Causes a warning to be logged when the number of routes defined by the <i>limit</i> argument have been redistributed. Additional redistribution is not prevented. |

Command Default

limit : 10240

threshold : 75 percent

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If someone mistakenly injects a large number of IPv6 routes into OSPFv3, perhaps by redistributing Border Gateway Protocol (BGP) into OSPFv3, the network can be severely flooded. Limiting the number of redistributed routes prevents this potential problem.

When the **maximum redistributed-prefixes** command is configured, if the number of redistributed routes reaches the maximum value configured, no more routes are redistributed (unless the **warning-only** keyword is configured).

The redistribution limit applies only to external IPv6 prefixes. Default routes and summarized routes are not limited.

The limit is tracked separately for each not-so-stubby-area (NSSA) because redistribution to NSSAs is done independently for each NSSA and independently of all other regular areas.

Select a maximum value based on your knowledge of how many prefixes are redistributed on the router to the OSPFv3 process.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to set a maximum of 2000 prefixes that can be redistributed into OSPFv3 process 1. If the number of prefixes redistributed reaches 75 percent of 2000 (1500 prefixes), a warning message is logged. Another warning is logged if the limit is reached and no more routes are redistributed.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute bgp 2406
RP/0/RSP0/CPU0:router(config-ospfv3)# maximum redistributed-prefixes 2000
```

mtu-ignore (OSPFv3)

To prevent the Open Shortest Path First Version 3 (OSPFv3) router process from checking whether neighbors are using the same maximum transmission unit (MTU) on a common interface when exchanging database descriptor (DBD) packets, use the **mtu-ignore** command in an appropriate configuration mode. To return to the default state, use the **no** form of this command.

mtu-ignore [disable]

no mtu-ignore

Syntax Description

| | |
|----------------|---|
| disable | (Optional) Disables the attribute in instances in which it is specified at a higher level in the configuration. |
| Note | The disable keyword is not available in router ospfv3 configuration mode. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the MTU ignore parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the MTU ignore parameter specified for the process.

If this command is not specified at any level, then OSPFv3 checks the MTU received from neighbors when exchanging DBD packets.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **mtu-ignore** command to check whether OSPFv3 neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the MTU configured on the incoming interface, OSPF adjacency is not established.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to disable MTU mismatch detection on received DBD packets on GigabitEthernet interface 0/1/0/3:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# mtu-ignore
```

neighbor (OSPFv3)

To configure Open Shortest Path First Version 3 (OSPFv3) routers interconnecting to nonbroadcast networks, use the **neighbor** command in interface configuration mode. To remove a configuration, use the **no** form of this command.

neighbor *ipv6-address* [**priority number**] [**poll-interval seconds**] [**cost number**] [**database-filter all out**]
no neighbor *ipv6-address* [**priority number**] [**poll-interval seconds**] [**cost number**] [**database-filter all out**]

Syntax Description

| | |
|-------------------------------------|---|
| <i>ipv6-address</i> | Link- local IP Version 6 (IPv6) address of the neighbor. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons. |
| priority <i>number</i> | (Optional) Specifies an 8-bit number indicating the router priority value of the nonbroadcast neighbor associated with the IP address specified. The priority keyword does not apply to point-to-multipoint interfaces. |
| poll-interval <i>seconds</i> | (Optional) Specifies an unsigned integer value (in seconds) reflecting the poll interval. RFC 1247 recommends that this value be much larger than the hello interval. The poll-interval keyword does not apply to point-to-multipoint interfaces. |
| cost <i>number</i> | (Optional) Assigns a cost to the neighbor, in the form of an integer from 1 to 65535. Neighbors with no specific cost configured assume the cost of the interface, based on the cost command. On point-to-multipoint interfaces, cost number is the only keyword and argument combination that works. The cost keyword does not apply to nonbroadcast multiaccess (NBMA) networks. |
| database-filter all out | (Optional) Filters outgoing link-state advertisements (LSAs) to an OSPFv3 neighbor. |

Command Default

No configuration is specified.

priority number : 0

poll-interval seconds : 120 seconds (2 minutes)

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

One neighbor entry must be included in the software configuration for each known nonbroadcast network neighbor. The neighbor address must be the IPv6 link-local address of the interface.

If a neighboring router has become inactive (hello packets have not been seen for the router dead interval period), it may still be necessary to send hello packets to the dead neighbor. These hello packets are sent at a reduced rate called the poll interval.

When the router starts up, it sends only hello packets to those routers with nonzero priority; that is, routers that are eligible to become designated routers (DRs) and backup designated routers (BDRs). After the DR and BDR are selected, the DR and BDR then start sending hello packets to all neighbors to form adjacencies.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to declare a router at address fe80::3203:a0ff:fe9d:f3fe on a nonbroadcast network:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# network non-broadcast
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# neighbor fe80::3203:a0ff:fe9d:f3fe
```

Related Commands

| Command | Description |
|---|---|
| priority (OSPFv3), on page 1249 | Sets the router priority, which helps determine the designated router for this network. |

network (OSPFv3)

To configure the Open Shortest Path First Version 3 (OSPFv3) network type to a type other than the default for a given medium, use the **network** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

```
network {broadcast| non-broadcast| {point-to-multipoint [non-broadcast]}| point-to-point}
no network
```

Syntax Description

| | |
|----------------------------|---|
| broadcast | Sets the network type to broadcast. |
| non-broadcast | Sets the network type to nonbroadcast multiaccess (NBMA). |
| point-to-multipoint | Sets the network type to point-to-multipoint. |
| [non-broadcast] | (Optional) Sets the point-to-multipoint network to be nonbroadcast. If you use the non-broadcast keyword, the neighbor command is required. |
| point-to-point | Sets the network type to point-to-point. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the network parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the network parameter specified for the process.

If this command is not specified at any level, then the OSPFv3 network type is the default of the given medium.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **network** command to configure broadcast networks as NBMA networks when, for example, routers in your network do not support multicast addressing.

Most times, it is assumed that when you configure NBMA networks as either broadcast or nonbroadcast, there are virtual circuits from every router to every router or fully meshed network. However, there are other configurations where this assumption is not true; for example, a partially meshed network. In these cases, you can configure the OSPFv3 network type as a point-to-multipoint network. Routing between two routers that are not directly connected goes through the router that has virtual circuits to both routers. You need not configure neighbors when using this command.

If the **network** command is issued on an interface that does not allow it, this command is ignored.

OSPFv3 has two features related to point-to-multipoint networks. One feature applies to broadcast networks and the other feature applies to nonbroadcast networks:

- On point-to-multipoint, broadcast networks, you can use the **neighbor** command, and you must specify a cost to that neighbor.
- On point-to-multipoint, nonbroadcast networks, you must use the **neighbor** command to identify neighbors. Assigning a cost to a neighbor is optional.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure an Ethernet interface as point-to-point:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface TenGigE0/1/0/3
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# network point-to-point
```

Related Commands

| Command | Description |
|---|---|
| neighbor (OSPFv3), on page 1236 | Configures OSPFv3 routers interconnecting to nonbroadcast networks. |

nssa (OSPFv3)

To configure an area as a not-so-stubby area (NSSA), use the **nssa** command in area configuration mode. To remove the NSSA distinction from the area, use the **no** form of this command.

nssa [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*| **metric-type** *type-value*]] [**no-summary**]

no nssa

Syntax Description

| | |
|--------------------------------------|---|
| no-redistribution | (Optional) Imports routes into the normal areas, but not into the NSSA area, by the redistribute command when the router is an NSSA area border router (ABR). |
| default-information-originate | (Optional) Generates a Type 7 default into the NSSA area. This keyword takes effect only on an NSSA ABR or NSSA autonomous system boundary router (ASBR). |
| metric <i>metric-value</i> | (Optional) Specifies a metric used for generating the default route. If you do not specify a default route metric value using the nssa and defaultmetric commands, the default metric value is 10. The value used is specific to the protocol. |
| metric-type <i>type-value</i> | (Optional) Specifies an external link type associated with the default route advertised into the Open Shortest Path First Version 3 (OSPFv3) routing domain. It can be one of the following values: 1 —Type 1 external route 2 —Type 2 external route |
| no-summary | (Optional) Prevents an (ABR) from sending summary link advertisements into the NSSA area. |

Command Default

No NSSA area is defined.

If you do not specify a value using the **default-metric** command, the default metric value is 10.

The default *type-value* is Type 2 external route.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A default route need not be defined in an NSSA ABR when the **nssa** command is configured. However, if this command is configured on an NSSA ASBR, then a default route must be defined.

**Note**

NSSA cannot be configured for area 0 (backbone area).

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure area 1 as an NSSA area:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.18.1.1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# nssa
```

nsr (OSPFv3)

To configure nonstop routing (NSR) for the Open Shortest Path First Version 3 (OSPFv3) protocol, use the **nsr** command in OSPFv3 router configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsr

no nsr

Syntax Description This command has no keywords or arguments.

Command Default NSR is not defined.

Command Modes OSPFv3 Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPFv3 NSR will be disabled at process startup, by default. When enabled, this state is remembered in the active process, and, is regardless of the presence and pairing state of a standby RP, as well as the state of the standby process.

NSR can be enabled for multiple OSPFv3 processes. The maximum number of processes on which NSR can be enabled is four.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure NSR for OSPFv3 process 211:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 211
RP/0/RSP0/CPU0:router(config-ospfv3)#nsr
```

Related Commands

| Command | Description |
|---|--------------------|
| router ospfv3, on page 1263 | |
| show ospfv3, on page 1267 | |

ospfv3 name-lookup

To configure Open Shortest Path First Version 3 (OSPFv3) to look up Domain Name System (DNS) names, use the **ospfv3 name-lookup** command in global configuration mode. To disable this function, use the **no** form of this command.

ospfv3 name-lookup

no ospfv3 name-lookup

Syntax Description This command has no arguments or keywords.

Command Default Routers are displayed by router ID or neighbor ID.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **ospfv3 name-lookup** command to simplify the task of searching for a router. Routers are displayed by name rather than by router ID or neighbor ID.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples The following example shows how to configure OSPFv3 to look up DNS names for use in all OSPFv3 **show** command displays:

```
RP/0/RSP0/CPU0:router(config)# ospfv3 name-lookup
```

packet-size (OSPFv3)

To configure the size of Open Shortest Path First Version 3 (OSPFv3) packets up to the size specified by the maximum transmission unit (MTU), use the **packet-size** command in an appropriate configuration mode. To disable this function and reestablish the default packet size, use the **no** form of this command.

packet-size *bytes*

no packet-size

Syntax Description

| | |
|--------------|---|
| <i>bytes</i> | Size in bytes. Range is 256 to 10000 bytes. |
|--------------|---|

Command Default

If not specified, the default packet size is 1500 bytes.

Command Modes

Router OSPFv3 configuration

Area configuration

Interface configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **packet-size** command to customize the size of OSPFv3 packets. The OSPFv3 protocol compares the packet size and the MTU size and uses the lower packet size value.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure the packet size:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# router ospf osp3  
RP/0/RSP0/CPU0:router(config-ospfv3)# packet-size 3500
```

passive (OSPFv3)

To suppress the sending of Open Shortest Path First Version 3 (OSPFv3) packets on an interface, use the **passive** command in an appropriate configuration mode. To remove the passive configuration, use the **no** form of this command.

passive [**disable**]

no passive

Syntax Description

| | |
|----------------|--|
| disable | (Optional) Sends OSPFv3 updates. |
| Note | The disable keyword is not available in router ospfv3 configuration mode. |

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the passive parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the passive parameter specified for the process.

If this command is not specified at any level, then the passive parameter is disabled and OSPFv3 updates are sent on the interface.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

OSPF routing information is neither sent nor received through the specified router interface. The specified interface address appears as a stub network in the OSPF domain.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| ospf | read, write |

Examples

The following example shows that OSPFv3 updates run over GigabitEthernet interface 0/3/0/0, 0/2/0/0, and 0/2/0/2. All other interfaces suppress sending OSPFv3 updates because they are in passive mode.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.0.0.206
RP/0/RSP0/CPU0:router(config-ospfv3)# passive
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# passive disable
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# exit
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# passive disable
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# passive
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/2
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# exit
```

priority (OSPFv3)

To set the router priority for an interface, which helps determine the designated router for an Open Shortest Path First Version 3 (OSPFv3) link, use the **priority** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

priority *value*

no priority

Syntax Description

| | |
|--------------|---|
| <i>value</i> | 8-bit unsigned integer indicating the router priority value. Range is 0 to 255. |
|--------------|---|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the priority parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the priority parameter specified by the process.

If this command is not specified at any level, then the default priority is 1.

Command Modes

Interface configuration

Area configuration

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When two routers attached to a network both attempt to become the designated router, the one with the higher router priority takes precedence. If there is a tie, the router with the higher router ID takes precedence. A router with a router priority set to zero is ineligible to become the designated router or backup designated router. Router priority is configured only for interfaces to broadcast and nonbroadcast multiaccess (NBMA) networks.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the router priority value to 4 on GigabitEthernet interface 0/1/0/1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/1/0/1
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# priority 4
```

Related Commands

| Command | Description |
|---|---|
| neighbor (OSPFv3), on page 1236 | Configures OSPFv3 routers interconnecting to nonbroadcast networks. |
| network (OSPFv3), on page 1238 | Configures the OSPFv3 network type to a type other than the default for a given medium. |

protocol shutdown (OSPFv3)

To disable an instance of the Open Shortest Path First protocol, version 3 (OSPFv3), so that it cannot form an adjacency on any interface, use the **protocol shutdown** command in the ospfv3 configuration mode. To re-enable the OSPF protocol, use the **no** form of this command.

protocol shutdown
no protocol shutdown

Command Default None

Command Modes ospfv3 configuration

| Command History | Release | Modification |
|-----------------|-------------|------------------------------|
| | Release 5.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **protocol shutdown** command to disable the OSPFv3 protocol for a specific routing instance without removing any existing OSPF configuration parameters.

The OSPFv3 protocol continues to run on the router and you can use the current OSPFv3 configuration, but OSPFv3 does not form any adjacencies on any interface.

This command is similar to performing the **no router ospf** command.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples This example shows how to disable the OSPFv3:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 ospfv3
RP/0/RSP0/CPU0:router(config-ospfv3)# protocol shutdown
```

range (OSPFv3)

To consolidate and summarize routes at an area boundary for Open Shortest Path First Version 3 (OSPFv3), use the **range** command in area configuration mode. To restore the default values, use the **no** form of this command.

range *ipv6-prefix/prefix-length* [**advertise**|**not-advertise**] [**cost** *number*]

no range *ipv6-prefix/prefix-length* [**advertise**|**not-advertise**] [**cost** *number*]

Syntax Description

| | |
|---------------------------|--|
| <i>ipv6-prefix</i> | Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons. |
| <i>/ prefix-length</i> | Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. |
| advertise | (Optional) Sets the address range status to advertise and generates a Type 3 summary link-state advertisement (LSA). |
| not-advertise | (Optional) Sets the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed and the component networks remain hidden from other networks. |
| cost <i>number</i> | (Optional) Specifies a cost for the range. Range is 1 to 16777214. |

Command Default

Routes are not consolidated and summarized for an area.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **range** command only with Area Border Routers (ABRs). It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing

information is condensed at area boundaries. External to the area, a single route is advertised for each address range. This process is called *route summarization*.

You can use the **range** command to configure multiple ranges. Thus, OSPFv3 can summarize addresses for many different sets of address ranges.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to specify one summary route to be advertised by the ABR to other areas for all IPv6 prefixes within the range defined by summary prefix 4004:f000::/32:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# range 4004:f000::/32
```

redistribute (OSPFv3)

To redistribute routes from one routing domain into Open Shortest Path First Version 3 (OSPFv3), use the **redistribute** command in an appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)

```
redistribute bgp process-id [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
no redistribute bgp process-id [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
```

Local Interface Routes

```
redistribute connected [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
no redistribute connected [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
```

Enhanced Interior Gateway Routing Protocol (EIGRP)

```
redistribute eigrp process-id [match {external [1|2]| internal}] [metric metric-value] [metric-type {1|2}] [route-policy policy-name] [tag tag-value]
no redistribute eigrp process-id [match {external [1|2]| internal}] [metric metric-value] [metric-type {1|2}] [route-policy policy-name] [tag tag-value]
```

Intermediate System-to-Intermediate System (IS-IS)

```
redistribute isis process-id [level-1| level-2| level-1-2] [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
no redistribute isis process-id [level-1| level-2| level-1-2] [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
```

Open Shortest Path First Version 3 (OSPFv3)

```
redistribute ospfv3 process-id match {external| 1| 2| internal| nssa-external| [1|2]} [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
no redistribute ospfv3 process-id [match {external| internal| nssa-external}] [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
```

Static

```
redistribute static [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
no redistribute static [metric metric-value] [metric-type {1|2}] [policy policy-name] [tag tag-value]
```

Syntax Description

| | |
|------------|---|
| bgp | Distributes routes from the BGP protocol. |
|------------|---|

| | |
|--|---|
| <i>process-id</i> | <p>For the bgp keyword, an autonomous system number has the following ranges:</p> <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. <p>For the isis keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> <p>For the ospf keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string.</p> |
| metric <i>metric-value</i> | (Optional) Specifies the metric used for the redistributed route. Range is 1 to 16777214. Use a value consistent with the destination protocol. |
| metric-type { 1 2 } | <p>(Optional) Specifies the external link type associated with the route advertised into the OSPF routing domain. It can be one of two values:</p> <ul style="list-style-type: none"> • 1—Type 1 external route • 2—Type 2 external route <p>If no metric-type is specified, the default is Type 2 external routes.</p> |
| tag <i>tag-value</i> | (Optional) Specifies the 32-bit dotted-decimal value attached to each external route. This value is not used by the OSPF protocol itself, but is carried in the External LSAs. Range is 0 to 4294967295. |
| policy <i>policy-name</i> | (Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to OSPF. |
| connected | Distributes routes that are established automatically by virtue of having enabled IP on an interface. |
| eigrp | Distributes routes from the EIGRP protocol. |
| isis | Distributes routes from the IS-IS protocol. |
| level-1 | (Optional) Redistributes Level 1 routes into other IP routing protocols independently. |
| level-1-2 | (Optional) Redistributes both Level 1 and Level 2 routes into other IP routing protocols. |
| level-2 | (Optional) Redistributes Level 2 routes into other IP routing protocols independently. |
| ospf | Distributes routes from the OSPF protocol. |

match { **internal** | **external** [**1** | **2**] | **nssa-external** [**1** | **2**] }

(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:

- **internal** —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes).
- **external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes.
- **nssa-external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes.

For the **external** and **nssa-external** options, if a type is not specified, then both Type 1 and Type 2 are assumed.

If no match is specified, the default is no filtering.

static Redistributes IP static routes.

Command Default

Route redistribution is disabled.

metric *metric-value*: Default is 20 for routes from all protocols except BGP routes, in which the default is 1.

metric-type *type-value* : Type 2 external route

All routes from the OSPFv3 routing protocol are redistributed.

tag *tag-value* : If no value is specified, the remote autonomous system number is used for routes from Border Gateway Protocol (BGP); for other protocols, the default is 0.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

**Note**

When redistributing routes (into OSPF) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Disabling or changing the arguments of any keyword does not affect the state of other keywords.

In general, route redistribution from Level 1 to Level 2 is automatic. You might want to use this command to better control which Level 1 routes can be redistributed into Level 2.

The redistribution of Level 2 routes into Level 1 is called *route leaking*. Route leaking is disabled by default. That is, Level 2 routes are not automatically included in Level 1 link-state protocols. If you want to leak Level 2 routes into Level 1, you must enable that behavior by using this command.

Redistribution from Level 1 into Level 1 and from Level 2 into Level 2 is not allowed.

A router receiving a link-state packet with an internal metric considers the cost of the route from itself to the redistributing router plus the advertised cost to reach the destination. An external metric considers only the advertised metric to reach the destination.

Redistributed routing information should always be filtered by the **distribute-list prefix-list out** command. Use of this command ensures that only those routes intended by the administrator are passed along to the receiving routing protocol.

OSPFv3 Considerations

Whenever you use the **redistribute** or the **default-information** command to redistribute routes into an OSPFv3 routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a default route into the OSPFv3 routing domain.

When routes are redistributed between OSPFv3 processes, no OSPFv3 metrics are preserved.

When routes are redistributed into OSPF and no metric is specified with the **metric** keyword, OSPF uses 20 as the default metric for routes from all protocols except BGP routes, which get a metric of 1. Furthermore, when the router redistributes from one OSPFv3 process to another OSPFv3 process on the same router, and if no default metric is specified, the metrics in one process are carried to the redistributing process.

BGP Considerations

The only connected routes affected by this command are the routes not specified by the **network** (BGP) command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to cause static routes to be redistributed into an OSPFv3 domain:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)# redistribute isis level-1
```

Related Commands

| Command | Description |
|--|--|
| default-information originate (OSPFv3), on page 1199 | Redistributes routes into a routing domain. |
| distribute-list prefix-list out, on page 1209 | Filters the routes redistributed into OSPFv3 from other routing protocols. |

retransmit-interval (OSPFv3)

To specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the Open Shortest Path First Version 3 (OSPFv3) interface, use the **retransmit-interval** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

retransmit-interval *seconds*

no retransmit-interval

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Time (in seconds) between retransmissions. It must be greater than the expected round-trip delay between any two routers on the attached network. Range is 1 to 65535 seconds. |
|----------------|--|

Command Default

If this command is not specified in interface configuration mode, then the interface adopts the retransmit interval parameter specified by the area.

If this command is not specified in area configuration mode, then the interface adopts the retransmit interval parameter specified by the process.

If this command is not specified at any level, then the default retransmit interval is 5 seconds.

Command Modes

Interface configuration
 Area configuration
 Router OSPFv3 configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA.

The setting of this parameter should be conservative, or needless retransmission results. The value should be larger for serial lines and virtual links.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the retransmit interval value to 8 seconds while in interface configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/2/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# retransmit-interval 8
```

router-id (OSPFv3)

To configure a router ID for the Open Shortest Path First Version 3 (OSPFv3) routing process, use the **router-id** command in an appropriate configuration mode. To cause the software to use the default method of determining the router ID, use the **no** form of this command after clearing or restarting the OSPF process.

router-id *router-id*

no router-id *router-id*

Syntax Description

| | |
|------------------|---|
| <i>router-id</i> | 32-bit router ID value specified in four-part, dotted-decimal notation. |
|------------------|---|

Command Default

If this command is not configured, the router ID is the highest IP address for an interface on the router, with any loopback interface taking precedence.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

We recommend that you use the **router-id** command to explicitly specify a unique 32-bit numeric value for the router ID. This configuration ensures that OSPFv3 can function regardless of the interface address configuration. Clear the OSPF process using the **clear ospf process** command or restart the OSPF process for the **no router-id** command to take effect.

In router OSPFv3 configuration mode, OSPF attempts to obtain a router ID in the following ways (in order of preference):

- 1 By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
- 2 The 32-bit numeric value specified by the OSPF **router-id** command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
- 3 A global router ID provided by the system (possibly, the first loopback address found at the boot time).

If the OSPFv3 process cannot obtain a router ID from any of these sources, the router issues the following error message:

```
%OSPFv3-4-NORTRID : OSPFv3 process 1 cannot run - configure a router ID for this process
At this point, OSPFv3 is effectively passive on all its interfaces. To run OSPFv3, make a router ID available by one of the methods described.
```

In VRF configuration mode, it is mandatory to configure a router ID manually. Otherwise, the OSPFv3 process will not become operational in the VRF.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to assign the IP address of 10.0.0.10 to the OSPFv3 process 109:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 109
RP/0/RSP0/CPU0:router(config-ospfv3)# router-id 10.0.0.10
```

Related Commands

| Command | Description |
|--|---|
| clear ospfv3 process, on page 1183 | Resets an OSPFv3 router process without stopping and restarting it. |

router ospfv3

To configure an Open Shortest Path First Version 3 (OSPFv3) routing process, use the **router ospfv3** command in global configuration mode. To terminate an OSPFv3 routing process, use the **no** form of this command.

router ospfv3 *process-name*

no router ospfv3 *process-name*

Syntax Description

| | |
|---------------------|--|
| <i>process-name</i> | Name that uniquely identifies an OSPFv3 routing process. The process name is any alphanumeric string no longer than 40 characters. |
|---------------------|--|

Command Default

No OSPFv3 routing process is defined.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can specify multiple OSPFv3 routing processes in each router. Up to 10 processes can be configured. The recommendation is not to exceed 4 OSPFv3 processes.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to instantiate an OSPFv3 routing process with a process name of 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
```

Related Commands

| Command | Description |
|--|--------------------------------------|
| area (OSPFv3), on page 1175 | Defines an OSPFv3 area. |
| interface (OSPFv3), on page 1224 | Defines an OSPFv3 interface by type. |

sham-link (OSPFv3)

To configure an Open Shortest Path First version 3 (OSPFv3) sham link between two provider edge routers on a non-default VRF, use the **sham-link** command in OSPFv3 router area sub-configuration mode. To terminate an OSPFv3 sham link, use the **no** form of this command.

sham-link *source-address destination-address*

no sham-link

Syntax Description

| | |
|----------------------------|--|
| <i>source-address</i> | IP address of the local (source) sham-link endpoint specified in four-part, dotted-decimal notation. |
| <i>destination-address</i> | IP address of the remote (destination) sham-link endpoint specified in four-part, dotted-decimal notation. |

Command Default

No sham link is configured.

Command Modes

OSPFv3 router area sub-configuration.

Command History

| Release | Modification |
|-------------|------------------------------|
| Release 5.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **sham-link** command to configure a point-to-point connection between two provider edge (PE) routers creating an interconnect between two VPN sites (VPN backbone). Sham links are configured on PE provider edge (PE) routers in a Multiprotocol Label Switching (MPLS) VPN backbone.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure an OSPFv3 sham link:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router ospfv3 ospfv3
```

```
RP/0/RSP0/CPU0:router(config_ospfv3)# vrf 1  
RP/0/RSP0/CPU0:router(config_ospfv3_vrf)# area 1  
RP/0/RSP0/CPU0:router(config_ospfv3_vrf_ar)# sham-link 100::1 200::1  
RP/0/RSP0/CPU0:router(config_ospfv3_vrf_ar_sl)# cost 23
```

Related Commands

| Command | Description |
|---|---|
| area (OSPFv3), on page 1175 | Configures an OSPF area. |
| cost (OSPFv3), on page 1191 | Explicitly specifies the cost of the interface (network) for OSPF path calculation. |

show ospfv3

To display general information about Open Shortest Path First Version 3 (OSPFv3) routing processes, use the **show ospfv3** command in EXEC mode.

```
show ospfv3 [process-name] [vrf {all | vrf-name} ]sham-links
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |
| sham-links | Displays OSPFv3 Sham-link information. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |
| Release 4.1.0 | The vrf <i>vrf_name</i> keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |
| Release 4.2.0 | Non-stop Routing (NSR) information was added in command output. |
| Release 5.1 | Displays OSPFv3 Sham-link information. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is a sample output from the **show ospfv3** command:

```
RP/0/RSP0/CPU0:router# show ospfv3 1

Routing Process "ospfv3 test" with ID 3.3.3.3
It is an autonomous system boundary router
Redistributing External Routes from,
static
  Maximum number of redistributed prefixes 10240
  Threshold for warning message 75%
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 255
Number of external LSA 1. Checksum Sum 0x004468
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Area BACKBONE(0) (Inactive)
    Number of interfaces in this area is 1
    SPF algorithm executed 1 times
    Number of LSA 3. Checksum Sum 0x018109
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

This table describes the significant fields shown in the display.

Table 118: show ospfv3 Field Descriptions

| Field | Description |
|--|---|
| Routing Process "ospfv3 test" with ID | OSPFv3 process name. |
| It is | Types are internal, area border, or autonomous system boundary. |
| Redistributing External Routes from | Lists of redistributed routes, by protocol. |
| Maximum number of redistributed prefixes | Number of redistributed prefixes. |
| Threshold for warning message | Warning message threshold. |
| Initial SPF schedule delay | Delay time of SPF calculations. |

| Field | Description |
|--|--|
| Minimum hold time between two consecutive SPFs | Minimum hold time between consecutive SPFs. |
| Maximum wait time between two consecutive SPFs | Maximum wait time between consecutive SPFs. |
| Initial LSA throttle delay | Delay time of LSA throttle. |
| Maximum hold time for LSA throttle | After initial throttle delay, the LSA generation is backed off by hold interval. |
| Maximum wait time for LSA throttle | Maximum throttle delay for LSA generation. |
| Minimum LSA arrival | Minimum LSA arrival. |
| LSA group pacing timer | Configured LSA group pacing timer (in seconds). |
| Interface flood pacing timer | Flooding pacing interval. |
| Retransmission pacing timer | Retransmission pacing interval. |
| Maximum number of configured interfaces | Maximum number of configured interfaces. |
| Number of external LSA | Number of external LSAs. |
| Number of areas in this router is | Number of areas configured for the router. |
| Number of interfaces in this area is | Number of interfaces in the area. |
| SPF algorithm executed <i>n</i> times | Times SPF algorithm was executed. |
| Number of LSA | Number of LSAs. |
| Number of DCbitless LSA | Number of DCbitless LSAs. |
| Number of indication LSA | Number of indication LSAs. |
| Number of DoNotAge LSA | Number of do-not-age LSAs. |
| Flood list length | Flood list length. |

This is sample output from the show ospfv3 vrf command that displays domain -id configuration:

```
RP/0/RSP0/CPU0:router#show ospfv3 0 vrf V1
Mon May 10 14:52:31.332 CEST

Routing Process "ospfv3 0" with ID 100.0.0.2 VRF V1
It is an area border and autonomous system boundary router
Redistributing External Routes from,
  bgp 1
  Maximum number of redistributed prefixes 10240
  Threshold for warning message 75%
Primary Domain ID:
```

show ospfv3

```

0x0005:0x00000000
Secondary Domain ID:
0x0105:0xbeef00000001
0x0205:0xbeef00000002
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 255
Maximum number of configured paths 16
Number of external LSA 2. Checksum Sum 0x015bb3
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Auto cost is enabled. Reference bandwidth 100
Area BACKBONE(0)
  Number of interfaces in this area is 1
  SPF algorithm executed 2 times
  Number of LSA 4. Checksum Sum 0x02629d
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0

```

This is sample output from the show ospfv3 vrf command that displays vrf-lite configuration:

```

RP/0/RSP0/CPU0:router#show ospfv3 0 vrf V2
Mon May 10 18:01:38.654 CEST

Routing Process "ospfv3 0" with ID 2.2.2.2 VRF V2
VRF lite capability is enabled
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 255
Maximum number of configured paths 16
Number of external LSA 0. Checksum Sum 00000000
Number of areas in this router is 0. 0 normal 0 stub 0 nssa
Auto cost is enabled. Reference bandwidth 100

```

This is sample output from the show ospfv3 command to verify that (Non-stop routing (NSR) is enabled:

```

RP/0/RSP0/CPU0:router#show ospfv3

Routing Process "ospfv3 100" with ID 3.3.3.3
NSR (Non-stop routing) is Enabled
It is an area border and autonomous system boundary router
Redistributing External Routes from,
  bgp 100
  Maximum number of redistributed prefixes 10240
  Threshold for warning message 75%
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Initial LSA throttle delay 0 msec
Minimum hold time for LSA throttle 5000 msec
Maximum wait time for LSA throttle 5000 msec
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec

```

```
Retransmission pacing timer 66 msec
Maximum number of configured interfaces 512
Maximum number of configured paths 16
Number of external LSA 0. Checksum Sum 00000000
Number of areas in this router is 15. 15 normal 0 stub 0 nssa
Auto cost is enabled. Reference bandwidth 100
```

The following is a sample output from the **show ospfv3** command with **sham-links** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 vrf vrf1 sham-links

Sham Links for OSPFv3 100, VRF vrf1

Sham Link OSPF_SL1 to address 300::1 is up
Area 2, source address 100::1
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:08
  Adjacency State FULL (Hello suppressed)
  Number of DBD retrains during last exchange 0
  Index 2/2, retransmission queue length 0, number of retransmission 0
  First 0(0)/0(0) Next 0(0)/0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
Sham Link OSPF_SL0 to address 200::1 is up
Area 2, source address 100::1
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:01
  Adjacency State FULL (Hello suppressed)
  Number of DBD retrains during last exchange 0
  Index 3/3, retransmission queue length 0, number of retransmission 0
  First 0(0)/0(0) Next 0(0)/0(0)
  Last retransmission scan length is 0, maximum is 0
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

Related Commands

| Command | Description |
|---|---------------------------------------|
| router ospfv3, on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 border-routers

To display the internal Open Shortest Path First Version 3 (OSPFv3) routing table entries to an area border router (ABR) and autonomous system boundary router (ASBR), use the **show ospfv3 border-routers** command in EXEC mode.

```
show ospfv3 [process-name] [vrf {all | vrf-name} ]border-routers [router-id]
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |
| <i>router-id</i> | (Optional) 32-bit router ID value specified in four-part, dotted-decimal notation. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf <i>vrf_name</i> keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospfv3 border-routers** command:

```
RP/0/RSP0/CPU0:router# show ospfv3 border-routers

OSPFv3 1 Internal Routing Table
Codes: i - Intra-area route, I - Inter-area route
i 10.0.0.207 [1] via fe80::3034:30ff:fe33:3742, GigabitEthernet 0/3/0/0, ABR/ASBR, Area
1, SPF 3
i 10.0.0.207 [10] via fe80::204:c0ff:fe22:73fe, Ethernet0/0/0/0, ABR/ASBR, Area 0, SPF 7
```

This table describes the significant fields shown in the display.

Table 119: show ospf border-routers Field Descriptions

| Field | Description |
|---------------------------|---|
| i | Type of this route; i indicates an intra-area route, I an inter-area route. |
| 10.0.0.207 | Router ID of destination. |
| [1] | Cost of using this route. |
| fe80::3034:30ff:fe33:3742 | Next-hop toward the destination. |
| GigabitEthernet 0/3/0/0 | Packets destined for fe80::3034:30ff:fe33:3742 are sent over GigabitEthernet interface 3/0/0/0. |
| ABR/ASBR | Router type of the destination; it is either an area border router (ABR) or autonomous system boundary router (ASBR) or both. |
| Area 1 | Area ID of the area from which this route is learned. |
| SPF 3 | Internal number of the shortest path first (SPF) calculation that installs this route. |

Related Commands

| Command | Description |
|--|---------------------------------------|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 database

To display lists of information related to the Open Shortest Path First Version 3 (OSPFv3) database for a specific router, use the **show ospfv3 database** command in EXEC mode.

```

show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database
show ospfv3 [process-name [area-id]] [vrf {all | vrf-name} ]database[adv-router [router-id]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [database-summary]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [external] [ link-state-id ]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [external] [ link-state-id ] [internal]
[adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [external] [ link-state-id ] [internal]
[self-originate]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [grace] [ link-state-id ] [adv-router
[ router-id ]] [internal] [self-originate]
show ospfv3 [process-name [area-id ]][vrf {all | vrf-name} ] database [inter-area prefix] [ link-state-id ]
show ospfv3 [process-name [area-id ]] vrf vrf-name database [vrf {all | vrf-name} ] [ link-state-id ] [internal]
[adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [inter-area prefix] [ link-state-id ]
[internal] [self-originate]
show ospfv3 [process-name [area-id ]][vrf {all | vrf-name} ] database [inter-area router] [ link-state-id ]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [inter-area router] [ link-state-id ]
[internal] [adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [inter-area router] [ link-state-id ]
[internal] [self-originate]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [link] [ link-state-id ]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [link] [ link-state-id ] [internal]
[adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [link] [ link-state-id ] [internal]
[self-originate]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [network] [ link-state-id ]
show ospfv3 [process-name [area-id ]][vrf {all | vrf-name} ] database [network] [ link-state-id ] [internal]
[adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] [vrf {all | vrf-name} ]database [network] [ link-state-id ] [internal]
[self-originate]
show ospfv3 [process-name [area-id ]][vrf {all | vrf-name} ] database [nssa-external] [ link-state-id ]
show ospfv3 [process-name [area-id ]][vrf {all | vrf-name} ] database [nssa-external] [ link-state-id ]
[internal] [adv-router [ router-id ]]
show ospfv3 [process-name [area-id ]] vrf vrf-name database [nssa-external] [ link-state-id ] [internal]
[self-originate]

```

```

show ospfv3 [process-name [ area-id ]][vrf {all | vrf-name} ] database [prefix] [ref-lsa] [router| network]
[ link-state-id ] [internal] [adv-router [ router-id ]]

show ospfv3 [process-name [ area-id ]][vrf {all | vrf-name} ] database [prefix] [ref-lsa] [router| network]
[ link-state-id ] [internal] [self-originate]

show ospfv3 [process-name [ area-id ]] vrf vrf-name database [prefix] [ link-state-id ]

show ospfv3 [process-name [ area-id ]][vrf {all | vrf-name} ] database [prefix] [ link-state-id ] [internal]
[adv-router [ router-id ]]

show ospfv3 [process-name [ area-id ]][vrf {all | vrf-name} ] database [prefix] [ link-state-id ] [internal]
[self-originate]

show ospfv3 [process-name [ area-id ]] [vrf {all | vrf-name} ]database [router] [ link-state-id ]

show ospfv3 [process-name area-id] [vrf {all | vrf-name} ]database [router] [adv-router [ router-id ]]

show ospfv3 [process-name [ area-id ]] [vrf {all | vrf-name} ]database [router] [ link-state-id ] [internal]
[self-originate]

show ospfv3 [process-name [ area-id ]][vrf {all | vrf-name} ] database [self-originate]

show ospfv3 [process-name [ area-id ]] [vrf {all | vrf-name} ]database [unknown [area| as| link]]
[ link-state-id ] [internal] [adv-router [ router-id ]] [self-originate]

```

Syntax Description

| | |
|--|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| adv-router [<i>router-id</i>] | (Optional) Displays all link-state advertisements (LSAs) of the specified router. |
| asbr-summary | (Optional) Displays information only about the Autonomous System Boundary Router (ASBR) summary LSAs. |
| database-summary | (Optional) Displays how many of each type of LSA are in the database for each area and the total. |
| external | (Optional) Displays information only about external LSAs. |
| grace | (Optional) Displays information about the state for the graceful restart link. |
| internal | (Optional) Displays information only about internal LSAs. |
| self-originate | (Optional) Displays only self-originated LSAs (from the local router). |
| <i>link-state-id</i> | (Optional) LSA ID that uniquely identifies the LSA. For network LSAs and link LSAs, this ID is the interface ID of the link of the router originating the LSA. |
| inter-area prefix | (Optional) Displays information only about the interarea prefix LSAs. |
| inter-area router | (Optional) Displays information only about the interarea router LSAs. |

| | |
|----------------------|---|
| link | (Optional) Displays information only about the link LSAs. |
| network | (Optional) Displays information only about the network LSAs. |
| nssa-external | (Optional) Displays information only about the not-so-stubby area (NSSA) external LSAs. |
| prefix | (Optional) Displays information only about the prefix LSAs. |
| ref-lsa | (Optional) Displays referenced LSA information. |
| router | (Optional) Displays information only about the router LSAs. |
| unknown | (Optional) Displays information only about unknown LSAs. |
| area | (Optional) Displays information only about the area LSAs. |
| as | (Optional) Displays information only about the autonomous system LSAs. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|---|
| | Release 3.7.2 | This command was introduced. |
| | Release 4.1.0 | The vrf <i>vrf_name</i> keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The various forms of this command deliver information about different OSPFv3 link-state advertisements.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospfv3 database** command with no arguments or keywords:

```
RP/0/RSP0/CPU0:router# show ospfv3 database
      OSPFv3 Router with ID (10.0.0.207) (Process ID 1)
      Router Link States (Area 0)
      ADV Router      Age          Seq#         Fragment ID  Link count  Bits
      0.0.0.1          163         0x80000039  0            2           None
      10.0.0.206       145         0x80000005  0            1           EB
      10.0.0.207       151         0x80000004  0            1           EB
      192.168.0.0      163         0x80000039  0            1           None

      Net Link States (Area 0)
      ADV Router      Age          Seq#         Link ID      Rtr count
      10.0.0.207      152         0x80000002  1            3
      192.168.0.0    163         0x80000039  1            2

      Inter Area Prefix Link States (Area 0)
      ADV Router      Age          Seq#         Prefix
      10.0.0.206      195         0x80000001  3002::/56
      10.0.0.207      197         0x80000001  3002::/56
      10.0.0.206      195         0x80000001  3002::206/128
      10.0.0.207      182         0x80000001  3002::206/128

      Inter Area Router Link States (Area 0)
      ADV Router      Age          Seq#         Link ID      Dest RtrID
      10.0.0.207      182         0x80000001  167772366   10.0.0.206
      10.0.0.206      182         0x80000001  167772367   10.0.0.207

      Link (Type-8) Link States (Area 0)
      ADV Router      Age          Seq#         Link ID      Interface
      0.0.0.1          163         0x80000039  1            Et0/0/0/0
      10.0.0.207      202         0x80000001  1            Et0/0/0/0
      10.0.0.206      200         0x80000001  2            Et0/0/0/0

      Intra Area Prefix Link States (Area 0)
      ADV Router      Age          Seq#         Link ID      Ref-1styp  Ref-LSID
      192.168.0.0    163         0x80000039  0            0x2002     1
      192.168.0.0    163         0x80000039  1            0x2001     0
      10.0.0.207     157         0x80000001  1001         0x2002     1
```

This table describes the significant fields shown in the display.

Table 120: show ospfv3 database Field Descriptions

| Field | Description |
|------------|---------------------------|
| ADV Router | ID of advertising router. |

| Field | Description |
|-------------|--|
| Age | Link-state age. |
| Seq# | Link-state sequence number (detects old or duplicate LSAs). |
| Fragment ID | Router LSA fragment ID. |
| Link count | Number of links described. |
| Bits | B indicates that the router is an area border router. E indicates that the router is an autonomous system boundary router. V indicates that the router is a virtual link endpoint. W indicates that the router is a wildcard multicast receiver. |
| Link ID | Unique LSA ID. |
| Rtr count | Number of routers attached to the link. |
| Prefix | Prefix of the route being described. |
| Dest RtrID | Router ID of the router being described. |
| Interface | Link described by the LSA. |
| Ref-lstype | LSA type of the LSA being referenced. |
| Ref-LSID | LSA ID of the LSA being referenced. |

The following is sample output from the **show ospfv3 database** command with the **external** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database external
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
          Type-5 AS External Link States
      LS age: 189
      LS Type: AS External Link
      Link State ID: 0
      Advertising Router: 10.0.0.206
      LS Seq Number: 80000002
      Checksum: 0xa303
      Length: 36
      Prefix Address: 2222::
      Prefix Length: 56, Options: None
      Metric Type: 2 (Larger than any link state path)
      Metric: 20
      External Route Tag: 0
```

This table describes the significant fields shown in the display.

Table 121: show ospfv3 database external Field Descriptions

| Field | Description |
|-----------------------|---|
| OSPFv3 Router with ID | Router ID number. |
| Process ID | OSPFv3 process name. |
| LS age | Link-state age. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID. |
| Advertising Router | ID of Advertising router. |
| LS Seq Number | Link-state sequence number (detects old or duplicate LSAs). |
| Checksum | LS checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Prefix Address | IPv6 address prefix of the route being described. |
| Prefix Length | Length of the IPv6 address prefix. |
| Metric Type | External type. |
| Metric | Link-state metric. |
| External Route Tag | External route tag, a 32-bit field attached to each external route. This tag is not used by the OSPFv3 protocol itself. |

The following is sample output from the **show ospfv3 database** command with the **inter-area prefix** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database inter-area prefix
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
            Inter Area Prefix Link States (Area 0)

LS age: 715
LS Type: Inter Area Prefix Links
Link State ID: 0
Advertising Router: 10.0.0.206
LS Seq Number: 80000002
Checksum: 0x3cb5
Length: 36
Metric: 1
Prefix Address: 3002::
```

Prefix Length: 56, Options: None

This table describes the significant fields shown in the display.

Table 122: show ospfv3 database inter-area prefix Field Descriptions

| Field | Description |
|-----------------------|--|
| OSPFv3 Router with ID | Router ID number. |
| Process ID | OSPFv3 process name. |
| LS age | Link-state age. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID. |
| Advertising Router | ID of advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Metric | Link-state metric. |
| Prefix Address | IPv6 prefix of the route being described. |
| Prefix Length | IPv6 prefix length of the route being described. |
| Options | LA indicates that the prefix is a local address. MC indicates the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at a not-so-stubby area (NSSA) area border. |

The following is sample output from the **show ospfv3 database** command with the **inter-area router** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database inter-area router
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
          Inter Area Router Link States (Area 0)

LS age: 1522
Options: (V6-Bit E-Bit R-bit DC-Bit)
LS Type: Inter Area Router Links
Link State ID: 167772366
Advertising Router: 10.0.0.207
```

```

LS Seq Number: 80000002
Checksum: 0xcaae
Length: 32
Metric: 1
Destination Router ID: 10.0.0.206

```

This table describes the significant fields shown in the display.

Table 123: show ospfv3 database inter-area routerField Descriptions

| Field | Description |
|-----------------------|---|
| OSPFv3 Router with ID | Router ID number. |
| Process ID | OSPFv3 process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA.) |
| Length | Length (in bytes) of the LSAs. |
| Metric | Link-state metric. |
| Destination Router ID | Router ID of the router being described. |

The following is sample output from the **show ospfv3 database** command with the **link** keyword:

```

RP/0/RSP0/CPU0:router# show ospfv3 database link
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
          Link (Type-8) Link States (Area 0)
LS age: 620
Options: (V6-Bit E-Bit R-bit DC-Bit)
LS Type: Link-LSA (Interface: Ethernet0/0/0/0)
Link State ID: 1 (Interface ID)

```

```

Advertising Router: 10.0.0.207
LS Seq Number: 80000003
Checksum: 0x7235
Length: 56
Router Priority: 1
Link Local Address: fe80::204:c0ff:fe22:73fe
Number of Prefixes: 1
Prefix Address: 7002::
Prefix Length: 56, Options: None

```

This table describes the significant fields shown in the display.

Table 124: show ospfv3 database link Field Descriptions

| Field | Description |
|---------------------------|---|
| OSPFv3 Router with ID | Router ID number. |
| Process ID | OSPFv3 process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports type-7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID (Interface ID). |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSAs. |
| Router Priority | Interface priority of originating router. |
| Link Local Address | Link local address of the interface. |
| Number of Prefixes | Number of prefixes associated with the link. |
| Prefix Address and Length | List of prefixes associated with the link. |

| Field | Description |
|---------|---|
| Options | LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates that the prefix should be propagated at an NSSA area border. |

The following is sample output from the **show ospfv3 database** command with the **network** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database network
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
          Net Link States (Area 0)
              LS age: 1915
              Options: (V6-Bit E-Bit R-bit DC-Bit)
              LS Type: Network Links
              Link State ID: 1 (Interface ID of Designated Router)
              Advertising Router: 10.0.0.207
              LS Seq Number: 80000004
              Checksum: 0x4330
              Length: 36
                  Attached Router: 10.0.0.207
                  Attached Router: 0.0.0.1
                  Attached Router: 10.0.0.206
```

This table describes the significant fields shown in the display.

Table 125: show ospfv3 database network Field Descriptions

| Field | Description |
|-----------------------|---|
| OSPFv3 Router with ID | Router ID number. |
| Process ID 1 | OSPFv3 process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID of the designated router. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |

| Field | Description |
|-----------------|--|
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Attached Router | List of routers attached to the network, by router ID. |

The following is sample output from the **show ospfv3 database** command with the **prefix** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database prefix
      OSPFv3 Router with ID (10.0.0.206) (Process ID 1)
          Intra Area Prefix Link States (Area 1)

Routing Bit Set on this LSA
LS age: 356
LS Type: Intra-Area-Prefix-LSA
Link State ID: 0
Advertising Router: 10.0.0.206
LS Seq Number: 8000001e
Checksum: 0xcdaa
Length: 44
Referenced LSA Type: 2001
Referenced Link State ID: 0
Referenced Advertising Router: 10.0.0.206
Number of Prefixes: 1
Prefix Address: 8006::
Prefix Length: 56, Options: None, Metric: 1
```

This table describes the significant fields shown in the display.

Table 126: show ospfv3 database prefix Field Descriptions

| Field | Description |
|-----------------------|--|
| OSPFv3 Router with ID | Router ID number. |
| Process ID 1 | OSPFv3 process name. |
| LS age | Link-state age. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID of the designated router. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |

| Field | Description |
|-------------------------------|--|
| Length | Length (in bytes) of the LSA. |
| Referenced LSA Type | Router LSA or network LSA of the prefixes referenced. |
| Referenced Link State ID | Link-state ID of the router or network LSA. |
| Referenced Advertising Router | Advertising router of the referenced LSA. |
| Number of Prefixes | Number of prefixes listed in the LSA. |
| Prefix Address | Prefix associated with the router or network. |
| Prefix Length | Length of the prefix. |
| Options | LA indicates that the prefix is a local address. MC indicates that the prefix is multicast capable. NU indicates that the prefix is not unicast capable. P indicates the prefix should be propagated at an NSSA area border. |
| Metric | Cost of the prefix. |

The following is sample output from the **show ospfv3 database** command with the **router** keyword:

```
RP/0/RSP0/CPU0:router# show ospfv3 database router
```

```

OSPFv3 Router with ID (10.0.0.206) (Process ID 1)

Router Link States (Area 0)

LS age: 814
Options: (V6-Bit E-Bit R-bit)
LS Type: Router Links
Link State ID: 0
Advertising Router: 0.0.0.1
LS Seq Number: 8000003c
Checksum: 0x51ca
Length: 56
Number of Links: 2

Link connected to: a Transit Network
Link Metric: 10
Local Interface ID: 1
Neighbor (DR) Interface ID: 1
Neighbor (DR) Router ID: 10.0.0.207

Link connected to: a Transit Network
Link Metric: 10
Local Interface ID: 2
Neighbor (DR) Interface ID: 1
Neighbor (DR) Router ID: 10.0.0.0

```

This table describes the significant fields shown in the display.

Table 127: show ospfv3 database router Field Descriptions

| Field | Description |
|-----------------------|---|
| OSPFv3 Router with ID | Router ID number. |
| Process ID 1 | OSPFv3 process name. |
| LS age | Link-state age. |
| Options | Type of service options (Type 0 only): DC—Supports demand circuits.E—Capable of processing external LSAs.MC—Forwards IP multicast.N—Supports Type 7 LSAs.R—Router is active.V6—Include in IPv6 routing calculations. |
| LS Type | Link-state type. |
| Link State ID | Link-state ID of the designated router. |
| Advertising Router | ID of the advertising router. |
| LS Seq Number | Link-state sequence (detects old or duplicate LSAs). |
| Checksum | Link-state checksum (Fletcher checksum of the complete contents of the LSA). |
| Length | Length (in bytes) of the LSA. |
| Link connected to | The type of network to which this interface is connected. Values are: <ul style="list-style-type: none"> • Another Router (point-to-point). • A Transit Network. • A Virtual Link. |
| Link Metric | OSPF cost of using this link. |
| Local Interface ID | Number that uniquely identifies an interface on a router. |

Related Commands

| Command | Description |
|--|---------------------------------------|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 flood-list

To display a list of Open Shortest Path First Version 3 (OSPFv3) link-state advertisements (LSAs) waiting to be flooded over an interface, use the **show ospfv3 flood-list** command in EXEC mode.

```
show ospfv3 [process-name] [area-id] [vrf {all | vrf-name} ]flood-list [type interface-path-id]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf_name keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 flood-list** command to display OSPFv3 packet pacing.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output from the **show ospfv3 flood-list** command shows three entries for the OSPFv3 1 process running over GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router# show ospfv3 flood-list GigabitEthernet 0/3/0/0
Flood Lists for OSPFv3 1
Interface GigabitEthernet 0/3/0/0, Queue length 3
Link state retransmission due in 24 msec
Displaying 3 entries from flood list:
Type      LS ID      ADV RTR    Seq NO      Age  Checksum
 3         0.0.0.199  10.0.0.207 0x80000002 3600 0x00c924
 3         0.0.0.200  10.0.0.207 0x80000002 3600 0x008966
 4         10.0.0.206 10.0.0.207 0x80000008    0 0x001951
```

This table describes the significant fields shown in the display.

Table 128: show ospfv3 flood-list Field Descriptions

| Field | Description |
|----------------------------------|---|
| Interface | Interface for which information is displayed. |
| Queue length | Number of LSAs waiting to be flooded. |
| Link state retransmission due in | Length of time before next link-state transmission. |
| Type | Type of LSA. |
| LS ID | Link-state ID of the LSA. |
| ADV RTR | IP address of advertising router. |
| Seq NO | Sequence number of LSA. |
| Age | Age of LSA (in seconds). |
| Checksum | Checksum of LSA. |

Related Commands

| Command | Description |
|--|---------------------------------------|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 interface

To display Open Shortest Path First Version 3 (OSPFv3) interface information, use the **show ospfv3 interface** command in EXEC mode.

```
show ospfv3 [ process-name ] [ area-id ] interface [vrf {all | vrf-name} ][type interface-path-id]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area number used to define the particular area. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf_name keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 interface** command when the adjacency between two neighboring routers is not forming. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the output to check the physical link and line protocol status and to confirm that the network type and timer intervals match those of the neighboring routers.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospfv3 interface** command when GigabitEthernet interface 0/2 /0/0 is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 interface GigabitEthernet 0/
2
/
0
/GigabitEthernet 0/2/0/0 is up, line protocol is up
Link Local address fe80::203:a0ff:fe9d:f3fe, Interface ID 2
Area 0, Process ID 1, Instance ID 0, Router ID 10.0.0.206
Network Type BROADCAST, Cost: 10
BFD enabled, interval 300 msec, multiplier 5
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 10.0.0.207, local address fe80::204:c0ff:fe22:73fe
Backup Designated router (ID) 10.0.0.206, local address fe80::203:a0ff:fe9d:f3fe
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
Index 0/2/1, flood queue length 0
Next 0(0)/0(0)/0(0)
Last flood scan length is 2, maximum is 9
Last flood scan time is 0 msec, maximum is 1 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Adjacent with neighbor 10.0.0.207 (Designated Router)
Suppress hello for 0 neighbor(s)
```

This table describes the significant fields shown in the display.

Table 129: show ospfv3 interface Field Descriptions

| Field | Description |
|--------------------|---|
| GigabitEthernet | Status of the physical link and operational status of the protocol. |
| Link Local Address | Interface link local address and interface ID. |

| Field | Description |
|----------------------------|--|
| Area | OSPFv3 area ID, process ID, instance ID, and router ID. |
| Transmit Delay | Transmit delay and interface state. |
| Designated Router | Designated router ID and respective interface IPv6 address. |
| Backup Designated router | Backup designated router ID and respective interface IPv6 address. |
| Timer intervals configured | Configuration of timer intervals. |
| Hello | Number of seconds until next hello packet is sent over this interface. |
| Index 0/2/1 | Link, area and autonomous system flood indexes, and number of flood queue entries. |
| Next 0(0)/0(0)/0(0) | Next link, area and autonomous system flood information, data pointer, and index. |
| Last flood scan length | Length of last flood scan. |
| Last flood scan time | Time of last flood scan (in milliseconds). |
| Neighbor Count | Count of network neighbors and list of adjacent neighbors. |
| Suppress hello | Count of neighbors suppressing hello messages. |

Related Commands

| Command | Description |
|--|---------------------------------------|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 message-queue

To display the information about the queue dispatch values, peak lengths, and limits, use the **show ospfv3 message-queue** command in EXEC mode.

```
show ospfv3 [process-name] [vrf {all | vrf-name} ] message-queue
```

Syntax Description

| | |
|-----------------|--|
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf <i>vrf_name</i> keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ospf | read |

Examples

This is sample output from the **show ospfv3 message-queue** command:

```
RP/0/RSP0/CPU0:router#show ospfv3 message-queue
Mon May 31 16:07:47.143 CEST

OSPFv3 Process 0
```

show ospfv3 message-queue

```
Hello Thread Packet Input Queue:
  Current queue length:      0
  Peak queue length:        2
  Queue limit:               5000
  Packets received:         104091
  Packets processed:        104091
  Packets dropped:          0
  Processing quantum:       10
  Full quantum used:        0
  Pulses sent:              104089
  Pulses received:         104089

Router Thread Message Queue
  Current queue length:      0
  Peak queue length:        2
  Low queue limit:          8000
  Medium queuing limit:     9000
  High queuing limit:       9500
  Messages queued:          1472
  Messages deleted:         0
  Messages processed:       1472
  Low queue drops:          0
  Medium queue drops:       0
  High queue drops:         0
  Processing quantum:       300
  Full quantum used:        0
  Pulses sent:              1484
  Pulses received:         1484
```

show ospfv3 neighbor

To display Open Shortest Path First Version 3 (OSPFv3) neighbor information on an individual interface basis, use the **show ospfv3 neighbor** command in EXEC mode.

```
show ospfv3 [process-name] [area-id] [vrf {all | vrf-name}] [neighbor [type interface-path-id] [neighbor-id] [detail]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) Neighbor router ID. |
| detail | (Optional) Displays all neighbors given in detail (lists all neighbors). |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf_name keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 neighbor** command when the adjacency between two neighboring routers is not forming. Adjacency means that the routers synchronize their databases when they discover each other.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following is sample output from the **show ospfv3 neighbor** command showing two lines of summary information for each neighbor:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor

Neighbors for OSPFv3 1

Neighbor ID      Pri   State           Dead Time   Interface ID  Interface
10.0.0.207      1     FULL/ -         00:00:35   3             GigabitEthernet 0/3/0/0

    Neighbor is up for 01:08:05
10.0.0.207      1     FULL/DR         00:00:35   2             Ethernet0/0/0/0
    Neighbor is up for 01:08:05

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 130: show ospfv3 neighbor Field Descriptions

| Field | Description |
|--------------|--|
| ID | Neighbor router ID. |
| Pri | Router priority for designated router election. A router with a priority of 0 is never elected as the designated router or backup designated router. |
| State | OSPFv3 state. |
| Dead Time | Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead. |
| Interface ID | Number that uniquely identifies an interface on a router. |
| Interface | Name of the interface that connects to this neighbor. |

| Field | Description |
|----------------|--|
| Neighbor is up | Time (in hh:mm:ss) that the OSPFv3 neighbor has been up. |

The following is sample output showing summary information about the neighbor that matches the neighbor ID:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor 10.0.0.207
```

```
Neighbors for OSPFv3 1

Neighbor 10.0.0.207
  In the area 0 via interface Ethernet0/0/0/0
  Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 10.0.0.207 BDR is 10.0.0.206
  Options is 0x13
  Dead timer due in 00:00:38
  Neighbor is up for 01:09:21
  Index 0/1/2, retransmission queue length 0, number of retransmission 1
  First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

Neighbor 10.0.0.207
  In the area 1 via interface GigabitEthernet 0/3/0/0
  Neighbor: interface-id 3, link-local address fe80::3034:30ff:fe33:3742
  Neighbor priority is 1, State is FULL, 6 state changes
  Options is 0x13
  Dead timer due in 00:00:38
  Neighbor is up for 01:09:21
  Index 0/1/1, retransmission queue length 0, number of retransmission 1
  First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

Total neighbor count: 2
```

This table describes the significant fields shown in the display.

Table 131: show ospfv3 neighbor 10.0.0.207 Field Descriptions

| Field | Description |
|--------------------|--|
| Neighbor | Neighbor router ID. |
| In the area | Area and interface through which the OSPFv3 neighbor is known. |
| link-local address | Link local address of the interface. |
| Neighbor priority | Router priority of neighbor and neighbor state. |
| State | OSPFv3 state. |
| state changes | Number of state changes for this neighbor. |

| Field | Description |
|----------------|--|
| DR is | Neighbor ID of the designated router. |
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub). |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead. |
| Neighbor is up | Time (in hh:mm:ss) that OSPFv3 neighbor has been up. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

The following sample output shows the neighbors that match the neighbor ID on the interface when the interface along with the neighbor ID is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1 10.0.0.207

Neighbors for OSPFv3 1

Neighbor 10.0.0.207
  In the area 0 via interface GigabitEthernet 0/3/0/1
  Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 10.0.0.207 BDR is 10.0.0.206
  Options is 0x13
  Dead timer due in 00:00:39
  Neighbor is up for 01:11:21
  Index 0/1/2, retransmission queue length 0, number of retransmission 1
  First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 132: show ospfv3 neighbor GigabitEthernet 0/3/0/1 10.0.0.207 Field Descriptions

| Field | Description |
|--------------------|--|
| Neighbor | Neighbor router ID. |
| In the area | Area and interface through which the OSPFv3 neighbor is known. |
| link-local address | Link local address of the interface. |

| Field | Description |
|-------------------|--|
| Neighbor priority | Router priority of neighbor and neighbor state. |
| State | OSPFv3 state. |
| state changes | Number of state changes for this neighbor. |
| DR is | Neighbor ID of the designated router. |
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub). |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead. |
| Neighbor is up | Time (in hh:mm:ss) that OSPFv3 neighbor has been up. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

The following sample output shows all neighbors on the interface when the interface is specified:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1

Neighbors for OSPFv3 1

Neighbor ID    Pri   State           Dead Time   Interface ID  Interface
10.0.0.207    1     FULL/DR         00:00:37   2             GigabitEthernet 0/3/0/1

Neighbor is up for 01:12:33

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 133: show ospfv3 neighbor GigabitEthernet 0/3/0/1 Field Descriptions

| Field | Description |
|-------------|--|
| Neighbor ID | Neighbor router ID. |
| Pri | Router priority for designated router election. A router with a priority of 0 is never elected as the designated router or backup designated router. |
| State | OSPF state. |

| Field | Description |
|----------------|--|
| Dead Time | Time (in hh:mm:ss) to elapse before OSPF declares the neighbor dead. |
| Interface ID | Number that uniquely identifies an interface on a router. |
| Interface | Name of the interface that connects to this neighbor. |
| Neighbor is up | Amount of time (in hh:mm:ss) that the OSPF neighbor has been up. |

The following is sample output showing detailed neighbor information for GigabitEthernet interface 0/3/0/1:

```
RP/0/RSP0/CPU0:router# show ospfv3 neighbor GigabitEthernet 0/3/0/1 detail

Neighbors for OSPFv3 1

Neighbor 10.0.0.207
  In the area 0 via interface GigabitEthernet 0/3/0/1
  Neighbor: interface-id 2, link-local address fe80::204:c0ff:fe22:73fe
  Neighbor priority is 1, State is FULL, 6 state changes
  DR is 10.0.0.207 BDR is 10.0.0.206
  Options is 0x13
  Dead timer due in 00:00:39
  Neighbor is up for 01:13:40
  Index 0/1/2, retransmission queue length 0, number of retransmission 1
  First 0(0)/0(0)/0(0) Next 0(0)/0(0)/0(0)
  Last retransmission scan length is 1, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec

Total neighbor count: 1
```

This table describes the significant fields shown in the display.

Table 134: show ospfv3 neighbor GigabitEthernet 0/3/0/1 detail Field Descriptions

| Field | Description |
|--------------------|--|
| Neighbor | Neighbor router ID. |
| In the area | Area and interface through which the OSPFv3 neighbor is known. |
| link-local address | Link local address of the interface. |
| Neighbor priority | Router priority of neighbor and neighbor state. |
| State | OSPFv3 state. |
| state changes | Number of state changes for this neighbor. |
| DR is | Neighbor ID of the designated router. |

| Field | Description |
|----------------|--|
| BDR is | Neighbor ID of the backup designated router. |
| Options | Hello packet options field contents (E-bit only; possible values are 0 and 2; 2 indicates that area is not a stub; 0 indicates that area is a stub). |
| Dead timer | Time (in hh:mm:ss) to elapse before OSPFv3 declares the neighbor dead. |
| Neighbor is up | Time (in hh:mm:ss) that the OSPFv3 neighbor has been up. |
| Index | Index and the remaining lines of this command give detailed information about flooding information received from the neighbor. |

Related Commands

| Command | Description |
|---|---------------------------------------|
| router ospfv3, on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 request-list

To display the first ten link-state requests pending that the local router is making to the specified Open Shortest Path First Version 3 (OSPFv3) neighbor and interface, use the **show ospfv3 request-list** command in EXEC mode.

```
show ospfv3 [process-name] [area-id] [vrf {all | vrf-name}] request-list [type interface-path-id] [neighbor-id]
```

Syntax Description

| | |
|--------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (Optional) Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) Router ID of the OSPFv3 neighbor. This argument must be in 32-bit dotted-decimal notation, similar to an IPv4 address. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf_name keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output shows request lists for neighbor 10.0.0.207 on the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 request-list 10.0.0.207 GigabitEthernet 0/3/0/0
Request Lists for OSPFv3 1
Neighbor 10.0.0.207, interface GigabitEthernet 0/3/0/0 address fe80::3034:30ff:fe33:3742

Type  LS ID          ADV RTR          Seq NO          Age  Checksum
  1    192.168.58.17    192.168.58.17   0x80000012     12  0x0036f3
  2    192.168.58.68    192.168.58.17   0x80000012     12  0x00083f
```

This table describes the significant fields shown in the display.

Table 135: show ospfv3 request-list Field Descriptions

| Field | Description |
|-----------|---|
| Neighbor | Router ID of the neighboring router. |
| interface | Name of the interface that connects to this neighbor. |
| address | IPv6 address of the neighbor. |
| Type | Type of link-state advertisement (LSA). |
| LS ID | Link-state ID of the LSA. |
| ADV RTR | Router ID of the advertising router. |

| Field | Description |
|----------|------------------------------|
| Seq NO | Sequence number of the LSA. |
| Age | Age of the LSA (in seconds). |
| Checksum | Checksum of the LSA. |

Related Commands

| Command | Description |
|---|--|
| router ospfv3, on page 1263 | Configures an OSPFv3 routing process. |
| show ospfv3 retransmission-list, on page 1305 | Displays the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface. |

show ospfv3 retransmission-list

To display the first ten link-state entries in the retransmission list that the local router sends to the specified neighbor over the specified interface, use the **show ospfv3 retransmission-list** command in EXEC mode.

```
show ospfv3 [process-name] [area-id] [vrf {all | vrf-name} ]retransmission-list [type interface-path-id]
[neighbor-id]
```

Syntax Description

| | |
|--------------------------|--|
| <i>process-name</i> | (Optional) Name that uniquely identifies an Open Shortest Path First Version 3 (OSPFv3) routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area-id</i> | (Optional) Area ID. If you do not specify an area, all areas are displayed. |
| <i>type</i> | (Optional) Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | (Optional) Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>neighbor-id</i> | (Optional) IP address of the OSPFv3 neighbor. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf vrf_name keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You might use this command when the databases of two neighboring routers are out of synchronization or if the adjacency is not forming between them. Adjacency means that the routers synchronize their databases when they discover each other.

You can look at the list to determine if one router is trying to request a particular database update. Entries that appear to be suspended in the list usually indicate that updates are not being delivered. One possible reason for this behavior is a maximum transmission unit (MTU) mismatch between the routers.

You might also look at this list to make sure it is not corrupted. The list should refer to database entries that actually exist.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output shows the retransmission list for neighbor 10.0.124.4 on GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router#show ospfv3 retransmission-list 10.0.124.4 GigabitEthernet 0/3/0/0
Neighbor 10.0.124.4, interface GigabitEthernet 0/3/0/0 address fe80::3034:30ff:fe33:3742
```

This table describes the significant fields shown in the display.

Table 136: show ospfv3 retransmission-list 10.0.124.4 GigabitEthernet 0/3/0/0 Field Descriptions

| Field | Description |
|-----------|---|
| Neighbor | Router ID of the neighboring router. |
| interface | Name of the interface that connects to this neighbor. |
| address | IPv6 address of the neighbor. |

Related Commands

| Command | Description |
|---|---------------------------------------|
| router ospfv3, on page 1263 | Configures an OSPFv3 routing process. |

| Command | Description |
|---|---|
| show ospfv3 request-list , on page 1302 | Displays the first ten link-state requests pending that the local router is making to the specified neighbor and interface. |

show ospfv3 routes

To display the Open Shortest Path First Version 3 (OSPFv3) route table, use the **show ospfv3 routes** command in EXEC mode.

show ospfv3 [*process-name*] **vrf** {**all** | *vrf-name*} **routes** [**external**| **connected**] [*ipv6-prefix/prefix-length*]

show ospfv3 [*process-name*] [**vrf** {**all** | *vrf-name*}]**routes summary**

Syntax Description

| | |
|------------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospf command. If this argument is included, only information for the specified routing process is displayed. |
| external | (Optional) Displays routes redistributed from other protocols. |
| connected | (Optional) Displays connected routes. |
| <i>ipv6-prefix</i> | (Optional) IP Version 6 (IPv6) prefix, which limits output to a specific route. This argument must be in the form documented in RFC 2373, in which the address is specified in hexadecimal using 16-bit values between colons. |
| <i>/ prefix-length</i> | (Optional) Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. |
| summary | Displays a summary of the route table. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| all | Display all the VRFs, excluding the default VRF. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

| Release | Modification |
|---------------|--|
| Release 4.1.0 | The <code>vrf vrf_name</code> keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 routes** command to display the OSPFv3 private routing table (which contains only routes calculated by OSPFv3). If there is something wrong with a route in the Routing Information Base (RIB), check the OSPFv3 copy of the route to determine if it matches the RIB contents. If it does not match, there is a synchronization problem between OSPFv3 and the RIB. If the routes match and the route is incorrect, OSPFv3 has made an error in its routing calculation.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output shows the route table for OSPFv3 process 1:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 routes
Route Table for OSPFv3 1 with ID 10.3.4.2
* 3000:11:22::/64, Inter, cost 21/0, area 1
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
  10.0.0.207/200
* 3000:11:22:1::/64, Inter, cost 31/0, area 1
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
  10.0.0.207/1
* 3333::/56, Ext2, cost 20/1, P:0 F:0
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
  10.0.0.207/0
* 6050::/56, Ext2, cost 20/1, P:0 F:0
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
  10.0.0.207/1
* 7002::/56, Intra, cost 10/0, area 0
  Ethernet0/0/0/0, connected
* 3000:11:22::/64, Inter, cost 21/0, area 1
  GigabitEthernet 0/3/0/0, fe80::3034:30ff:fe33:3742
  10.0.0.207/200
```

This table describes the significant fields shown in the display.

Table 137: show ospfv3 1 route Field Descriptions

| Field | Description |
|-----------------|-----------------------------------|
| 3000:11:22::/64 | Route prefix to the local router. |

show ospfv3 routes

| Field | Description |
|---------------------------|--|
| Inter | Prefix 3000:11:22::/64 is interarea. |
| cost 21/0 | Sum of the link costs required to reach prefix 3000:11:22::/64. 0. In this example, 20 is the external cost. |
| GigabitEthernet 0/3/0/0 | Packets destined for prefix 3000:11:22::/64 are sent over the GigabitEthernet 0/3/0/0 interface. |
| fe80::3034:30ff:fe33:3742 | Next-hop router on the path to prefix 3000:11:22::/64. |
| 10.0.0.207 | Router 10.0.0.207 is the router that advertised this route. |

Related Commands

| Command | Description |
|---|---------------------------------------|
| router ospfv3, on page 1263 | Configures an OSPFv3 routing process. |

show ospfv3 statistics rib-thread

To display RIB thread statistics, use the **show ospfv3 statistics rib-thread** command in EXEC mode.

```
show ospfv3 [process-name [area-id]] statistics rib-thread
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPF routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| <i>area id</i> | (Optional) Area number used to define the particular area. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ospf | read |

Examples

This is sample output from **show ospfv3 statistics rib-thread** command:

```
RP/0/RSP0/CPU0:router#show ospfv3 0 statistics rib-thread
Mon May 10 17:48:29.011 CEST
OSPFv3 0 RIB thread statistics
Queue statistics:
  Last entry dequeue          10127056 msecs ago (14:59:42.171)
  RIB thread active           NO
  Total RIB thread signals    30
  Current queue length        0
  Maximum queue length        2
  Total entries queued         31
  Total entries dequeued       31
```

show ospfv3 statistics rib-thread

```
Maximum latency (msec)      5.000
Average latency (msec)      0.323
Queue errors:
  Enqueue errors            0
  Dequeue errors            0
RIB batch statistics:
  Batches sent to RIB       31
  Batch all routes OK       31
  Batch some routes backup  0
RIB batch errors:
  Batches version mismatch  0
  Batches missing connection 0
  Batches no table          0
  Batch route table limit   0
  Batch route errors        0
  Batch errors              0
  Route table limit         0
  Route path errors         0
  Route errors              0
  Path table limit          0
  Path errors               0
```

show ospfv3 summary-prefix

To display Open Shortest Path First Version 3 (OSPFv3) aggregated summary address information, use the **show ospfv3 summary-prefix** command in EXEC mode.

```
show ospfv3 [process-name] [vrf vrf-name] summary-prefix
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf { <i>vrf_name</i> } keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show ospfv3 summary-prefix** command if you configured summarization of external routes with the **summary-prefix** command and you want to display configured summary addresses.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output shows the summary prefix address for the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 summary-prefix
OSPFv3 Process 1, Summary-prefix
4004:f000::/32 Metric 20, Type 2, Tag 0
```

This table describes the significant fields shown in the display.

Table 138: show ospfv3 1 summary-prefix Field Descriptions

| Field | Description |
|----------------|--|
| 4004:f000::/32 | Summary prefix designated for a range of IPv6 prefixes. The length of the IPv6 prefix. |
| Metric | Metric used to advertise the summary routes. |
| Type | External link-state advertisements (LSAs) metric type. |
| Tag | Tag value that can be used as a “match” value for controlling redistribution through route maps. |

Related Commands

| Command | Description |
|--|---|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |
| summary-prefix (OSPFv3) , on page 1331 | Creates aggregate addresses for routes being redistributed from another routing protocol into OSPFv3. |

show ospfv3 virtual-links

To display parameters and the current state of Open Shortest Path First Version 3 (OSPFv3) virtual links, use the **show ospfv3 virtual-links** command in EXEC mode.

```
show ospfv3 [process-name] [vrf vrf-name] virtual-links
```

Syntax Description

| | |
|---------------------|---|
| <i>process-name</i> | (Optional) Name that uniquely identifies an OSPFv3 routing process. The process name is defined by the router ospfv3 command. If this argument is included, only information for the specified routing process is displayed. |
| vrf | Specifies an OSPF VPN routing and forwarding (VRF) instance. |
| <i>vrf-name</i> | Specifies the name of the of the OSPFv3 VRF. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | The vrf { <i>vrf_name</i> } keyword and argument were added to show OSPFv3 VRF configuration information. The output of the command was enhanced to include the name of the VRF. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The information displayed in the **show ospfv3 virtual-links** command is useful in debugging OSPFv3 routing operations.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following sample output shows the virtual links for the OSPFv3 1 process:

```
RP/0/RSP0/CPU0:router# show ospfv3 1 virtual-links

Virtual Links for OSPFv3 1
Virtual Link to router 172.31.101.2 is up
Interface ID 16, IPv6 address 3002::206
Transit area 0.0.0.1, via interface GigabitEthernet 0/3/0/0, Cost of using 11
Transmit Delay is 5 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 0:00:08
Adjacency State FULL
```

This table describes the significant fields shown in the display.

Table 139: show ospfv3 virtual-links Field Descriptions

| Field | Description |
|------------------------------|--|
| Virtual Link to router is up | Specifies the OSPFv3 neighbor, and if the link to that neighbor is up or down. |
| Interface ID | ID of the virtual link interface. |
| IPv6 address | IPv6 address of virtual link endpoint. |
| Transit area | Transit area through which the virtual link is formed. |
| via interface | Interface through which the virtual link is formed. |
| Cost | Cost of reaching the OSPF neighbor through the virtual link. |
| Transmit Delay | Transmit delay on the virtual link. |
| State POINT_TO_POINT | State of the OSPFv3 neighbor. |
| Timer intervals | Various timer intervals configured for the link. |
| Hello due in | When the next hello message is expected from the neighbor (in hh:mm:ss). |
| Adjacency State | Adjacency state between the neighbors. |

Related Commands

| Command | Description |
|--|---------------------------------------|
| router ospfv3 , on page 1263 | Configures an OSPFv3 routing process. |

show protocols (OSPFv3)

To display information about the Open Shortest Path First Version 3 (OSPFv3) process running on the router, use the **show protocols** command in EXEC mode.

```
show protocols [afi-all| ipv4| ipv6] [all| protocol]
```

Syntax Description

| | |
|-----------------|---|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies an IPv4 address family. |
| ipv6 | (Optional) Specifies an IPv6 address family. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. For the IPv4 address family, the options are: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip <p>For the IPv6 address family, the options are:</p> <ul style="list-style-type: none"> • bgp • eigrp • isis • ospfv3 |

Command Default

The default address family is IPv4.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

| Release | Modification |
|---------------|--|
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ospf | read |

Examples

The following example is sample output from the **show protocols** command:

```
RP/0/RSP0/CPU0:router# show protocols ipv6 ospfv3

Routing Protocol OSPFv3 1
  Router Id:10.0.0.1
  Distance:110
  Redistribution:
    None
  Area 0
    GigabitEthernet 0/2/0/2
    Loopback1
```

This table describes the significant fields shown in the display.

Table 140: show protocols Field Descriptions

| Field | Description |
|----------------|--|
| Router Id | Router ID of the OSPFv3 process. |
| Distance | Administrative distance for the protocol. This distance determines the priority the Routing Information Base (RIB) gives to the routes, as opposed to other protocols, for example, IS-IS. |
| Redistribution | Protocols from which this OSPFv3 process is redistributing routes. |
| Area | OSPFv3 areas defined in this process, followed by their associated interfaces. |

snmp context (OSPFv3)

To specify an SNMP context for an OSPFv3 instance, use the **snmp context** command in router configuration mode or in VRF configuration mode. To remove the SNMP context, use the **no** form of this command.

snmp context *context_name*

no snmp context *context_name*

Syntax Description

| | |
|---------------------|---|
| <i>context_name</i> | Specifies name of the SNMP context for OSPFv3 instance. |
|---------------------|---|

Command Default

SNMP context is not specified.

Command Modes

Router OSPFv3 configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes the proper task IDs. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The `snmp-server` commands need to be configured to perform SNMP request for the OSPF instance. Refer *SNMP Server Commands* module in *Cisco ASR 9000 Series Aggregation Services Router System Management Command Reference* for information on using the `snmp-server` commands.



Note

To map an SNMP context with a protocol instance, topology or VRF entity, use the **snmp-server context mapping** command. However, the **feature** option of this command does not work with OSPFv3 protocol.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure an SNMP context *foo* for OSPFv3 instance *100*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 100
RP/0/RSP0/CPU0:router(config-ospf)#snmp context foo
```

This example shows how to configure **snmp-server** commands to be used with the **snmp context** command:

```
RP/0/RSP0/CPU0:router(config)#snmp-server host 10.0.0.2 traps version 2c public udp-port
1620
RP/0/RSP0/CPU0:router(config)#snmp-server community public RW
RP/0/RSP0/CPU0:router(config)#snmp-server contact foo
RP/0/RSP0/CPU0:router(config)#snmp-server community-map public context foo
```

This is a sample SNMP context configuration for OSPFv3 instance *100*:

```
snmp-server host 10.0.0.2 traps version 2c public udp-port 1620
snmp-server community public RW
snmp-server contact foo

snmp-server community-map public context foo

router ospfv3 100
  router-id 2.2.2.2
  bfd fast-detect
  nsf cisco
  snmp context foo
  area 0
    interface Loopback1
    !
  !
  area 1
    interface GigabitEthernet0/2/0/1
      demand-circuit enable
    !
    interface POS0/3/0/0
    !
    interface POS0/3/0/1
    !
  !
  !
  !
```

Related Commands

| Command | Description |
|--|---|
| snmp trap (OSPFv3), on page 1322 | Enables SNMP trap for an OSPFv3 instance. |
| snmp-server host | Specifies the recipient of an SNMP notification operation. |
| snmp-server community | Configures the community access string to permit access to the Simple Network Management Protocol (SNMP). |
| snmp-server contact | Sets the Simple Network Management Protocol (SNMP) system contact. |

| Command | Description |
|----------------------------------|--|
| snmp-server community-map | Associates a Simple Network Management Protocol (SNMP) community with an SNMP context. |

snmp trap (OSPFv3)

To enable SNMP trap for an OSPFv3 instance, use the **snmp trap** command in VRF configuration mode. To disable SNMP trap for the OSPFv3 instance, use the **no** form of this command.

snmp trap

no snmp trap

Syntax Description This command has no keywords or arguments.

Command Default Disabled.

Command Modes OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to enable SNMP trap for OSPFv3 instance *100* under VRF *vrf-1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 100
RP/0/RSP0/CPU0:router(config-ospf)#vrf vrf-1
RP/0/RSP0/CPU0:router(config-ospf-vrf)#snmp trap
```

Related Commands

| Command | Description |
|---|--|
| snmp context (OSPFv3), on page 1319 | Specifies SNMP context for an OSPFv3 instance. |

snmp trap rate-limit (OSPFv3)

To control the number of traps that OSPFv3 sends by configuring window size and the maximum number of traps during that window, use the **snmp trap rate-limit** command in router OSPFv3 configuration mode or OSPFv3 VRF configuration mode. To disable configuring the window size and maximum number of traps during the window, use the **no** form of this command.

snmp trap rate-limit *window-size max-num-traps*

no snmp trap rate-limit *window-size max-num-traps*

Syntax Description

| | |
|----------------------|---|
| <i>window-size</i> | Specifies the trap rate limit sliding window size. The range is 2 to 60 windows. |
| <i>max-num-traps</i> | Specifies the maximum number of traps sent in window time. The range is 0 to 300 traps. |

Command Default

None

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to set the trap rate limit sliding window size to 50 and the maximum number of traps sent to 250 for OSPFv3 instance *100* under vrf *vrf1*:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#router ospfv3 100  
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrfl  
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#snmp trap rate-limit 50 250
```

spf prefix-priority (OSPFv3)

To prioritize OSPFv3 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run, use the **spf prefix-priority** command in router configuration mode or VRF configuration mode. To return to the system default value, use the **no** form of this command.

spf prefix-priority route-policy *policy-name* [**disable**]

spf prefix-priority route-policy *policy-name*

Syntax Description

| | |
|---------------------|--|
| route-policy | Specifies the route-policy to prioritize route installation. |
| <i>policy-name</i> | Name of the route policy. |
| disable | Disables SPF prefix priority |

Command Default

SPF prefix prioritization is disabled.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure OSPFv3 SPF prefix prioritization:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# prefix-set ospf3-critical-prefixes
RP/0/RP0/CPU0:router(config-pfx)# 66.0.0.0/16
RP/0/RP0/CPU0:router(config-pfx)# end-set
```

```

RP/0/RP0/CPU0:router(config)# route-policy ospf3-spf-priority
RP/0/RP0/CPU0:router(config-rpl)# if destination in ospf-critical-prefixes then set
spf-priority critical
endif
RP/0/RP0/CPU0:router(config-rpl)# end-policy
RP/0/RP0/CPU0:router(config-rpl)# commit
RP/0/RP0/CPU0:router(config-rpl)# exit
RP/0/RP0/CPU0:router(config)# router ospfv3 1
RP/0/RP0/CPU0:router(config-ospf)# router-id 66.0.0.1
RP/0/RP0/CPU0:router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority

```

Related Commands

| Command | Description |
|---------------------------|--|
| prefix-set | Enters prefix set configuration mode and defines a prefix set. |
| route-policy (RPL) | Defines a route policy and enters route-policy configuration mode. |

stub (OSPFv3)

To define an area as a stub area for Open Shortest Path First Version 3 (OSPFv3), use the **stub** command in area configuration mode. To disable this function, use the **no** form of this command.

stub [no-summary]

no stub

Syntax Description

| | |
|-------------------|---|
| no-summary | (Optional) Prevents an area border router (ABR) from sending summary link advertisements into the stub area. Areas with this option are known as <i>totally stubby</i> areas. |
|-------------------|---|

Command Default

No stub area is defined.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You must configure the **stub** command on all routers in the stub area. Use the **default-cost area** command on the ABR of a stub area to specify the cost of the default route advertised into the stub area by the ABR.

Two stub area router configuration commands exist: the **stub** and **default-cost** commands. In all routers attached to the stub area, the area should be configured as a stub area using the **stub** command. Use the **default-cost** command only on an ABR attached to the stub area. The **default-cost** command provides the metric for the summary default route generated by the ABR into the stub area.

To further reduce the number of link-state advertisements (LSAs) sent into a stub area, you can configure the **no-summary** keyword on the ABR to prevent it from sending summary LSAs (LSA Type 3) into the stub area.

A stub area does not accept information about routes external to the autonomous system.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to create stub area 5 and specifies a cost of 20 for the default summary route sent into this stub area:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 5
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# stub
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# default-cost 20
```

Related Commands

| Command | Description |
|---|---|
| default-cost (OSPFv3), on page 1197 | Specifies a cost for the default summary route sent into a stub area. |

stub-router

To modify self originated router LSAs when stub router is active, use the **stub-router** command in an appropriate configuration mode. To disable this function, use the **no** form of this command.

```
stub-router router-lsa [r-bit| v6-bit| max-metric] [always] [on-proc-migration interval] [on-proc-restart interval] [on-switchover interval] [on-startup [interval] wait-for-bgp] [summary-lsa [metric ]] [external-lsa [metric ]] [include-stub]
```

```
stub-router router-lsa [r-bit| v6-bit| max-metric]
```

Syntax Description

| | |
|--------------------------|---|
| router-lsa | Specifies that always originate router link-state advertisements (LSAs) with the stub-router. |
| r-bit | Router-LSAs are originated with R-bit clear (v6 bit set), which means the node does not act as a transit router. Directly connected networks (native to OSPF) are still reachable within the OSPF area. |
| v6-bit | Router-LSAs are originated with V6 bit clear (and also r-bit clear). That means the node is not willing to receive any ipv6 traffic. Other ospfv3 routers won't install any route to a node with v6-bit clear. |
| max-metric | Router-LSAs are originated with maximum metric. Unlike the r-bit and v6-bit mode, the router may still act as a transit node, if there is no alternate path. |
| always | Stub-router mode is activated unconditionally. |
| on-proc-migration | Stub-router mode is activated for the desired period of time, upon ospfv3 process migration. |
| on-proc-restart | Stub-router mode is activated for the desired period of time, upon ospfv3 process restart. |
| on-switchover | Stub-router mode is activated for the desired period of time, upon RP failover. |
| on-startup | Stub-router mode is activated (for configured time, or until BGP converges) upon router startup (boot). |
| wait-for-bgp | Stub-router mode is terminated upon BGP convergence in ipv6 unicast address family. This option could only be used in the global routing table, not in a non-default VRF. This option is only supported with the on-startup trigger when the router boots. |
| summary-lsa | <p>If enabled, summary LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode.</p> <p>In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear).</p> <p>If enabled and metric is not explicitly configured, the default metric for summary LSAs when stub-router active is 16711680 (0xFF0000).</p> |

| | |
|---------------------|---|
| external-lsa | <p>If enabled, external LSAs are advertised with modified metric when stub-router is active. This configuration is applicable to max-metric mode.</p> <p>In r-bit mode, ABR/ASBR functionality is implicitly disabled and routers will not use this node as an ABR/ASBR, since it declares no transit capability (r-bit clear).</p> <p>If enabled and metric is not explicitly configured, the default metric for external LSAs when stub-router active is 16711680 (0xFF0000).</p> |
| include-stub | <p>If enabled, intra-area-prefix LSAs that are referencing router LSA are advertised with maximum metric (0xffff) when stub-router is active.</p> <p>Intra-area-prefix LSAs that are referencing network LSA do not change metric</p> <p>Can be used in r-bit and max-metric modes.</p> <p>/128 prefixes that are normally advertised with LA-bit set and 0 metric are also advertised with maximum metric and LA-bit clear when stub-router is active.</p> |

Command Default Disabled.

Command Modes Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Only one method (r-bit, v6-bit, max-metric) could be activated at a time. Configuring the methods simultaneously, or different method per trigger, is not supported.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure router LSAs are originated with R-bit clear under OSPFv3 VRF, *vrf_1*:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#stub-router router-lsa r-bit
```

summary-prefix (OSPFv3)

To create aggregate addresses for routes being redistributed from another routing protocol into Open Shortest Path First Version 3 (OSPFv3) protocol, use the **summary-prefix** command in an appropriate configuration mode. To stop summarizing redistributed routes, use the **no** form of the command.

summary-prefix *ipv6-prefix/prefix-length* [**not-advertise**] **tag** *tag*

no summary-prefix *ipv6-prefix/prefix-length*

Syntax Description

| | |
|------------------------|---|
| <i>ipv6-prefix</i> | Summary prefix designated for a range of IP Version 6 (IPv6) prefixes. This argument must be in the form documented in RFC 2373, where the address is specified in hexadecimal using 16-bit values between colons. |
| <i>/ prefix-length</i> | Length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. |
| not-advertise | (Optional) Suppresses summary routes that match the address and mask pair from being advertised. |
| tag tag | (Optional) Specifies a tag value that can be used as a “match” value for controlling redistribution. |

Command Default

When this command is not used in router configuration mode, aggregate addresses are not created for routes being redistributed from another routing protocol into the OSPFv3 protocol.

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **summary-prefix** command to cause an OSPFv3 Autonomous System Boundary Router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. This command summarizes only routes from other routing protocols that are being redistributed into OSPFv3.

You can use this command multiple times to summarize multiple groups of addresses. The metric used to advertise the summary is the lowest metric of all the more specific routes. This command helps reduce the size of the routing table.

If you want to summarize routes between OSPFv3 areas, use the **range** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

In the following example, if summary prefix 4004:f000:132 is configured and routes 4004:f000:1::/64, 4004:f000:2::/64, and 4004:f000:3::/64 are redistributed into OSPFv3; only route 4004:f000::/32 is advertised in an external link-state advertisement:

```
RP/0/RSP0/CPU0:router(config-ospfv3)# summary-prefix 4004:f000::/32
```

Related Commands

| Command | Description |
|--|---|
| range (OSPFv3), on page 1252 | Consolidates and summarizes routes at an area boundary. |

timers lsa arrival

To set the minimum interval at which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First Version 3 (OSPFv3) neighbors, use the **timers lsa arrival** command in an appropriate configuration mode. To restore the default value, use the **no** form of this command.

timers lsa arrival *milliseconds*

no timers lsa arrival

Syntax Description

| | |
|---------------------|--|
| <i>milliseconds</i> | Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. Range is 0 to 60000 milliseconds. |
|---------------------|--|

Command Default

1000 milliseconds

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers lsa arrival** command to control the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped.

We recommend that the *milliseconds* value of the **timers lsa arrival** command be less than or equal to the *hold-interval* value of the **timers throttle lsa all** command for the neighbor.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1  
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000  
RP/0/RSP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

Related Commands

| Command | Description |
|--|---|
| timers throttle lsa all (OSPFv3), on page 1341 | Sets rate-limiting values for LSAs being generated. |

timers pacing flood

To configure link-state advertisement (LSA) flood packet pacing, use the **timers pacing flood** command in an appropriate configuration mode. To restore the default flood packet pacing value, use the **no** form of this command.

timers pacing flood *milliseconds*

no timers pacing flood

Syntax Description

milliseconds

Time (in milliseconds) at which LSAs in the flooding queue are paced in between updates. Range is 5 milliseconds to 100 milliseconds.

Command Default

milliseconds: 33

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configuring OSPFv3 flood pacing timers allows you to control interpacket spacing between consecutive link-state update packets in the OSPF transmission queue. Use the **timers pacing flood** command to control the rate at which LSA updates occur, thereby preventing high CPU or buffer utilization that can result when an area is flooded with a very large number of LSAs.

The default settings for OSPFv3 packet pacing timers are suitable for the majority of OSPFv3 deployments. Do not change the packet pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flood timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes risks associated with changing the default flood timer values.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure LSA flood packet-pacing updates to occur in 55-millisecond intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing flood 55
```

Related Commands

| Command | Description |
|---|--|
| show ospfv3, on page 1267 | Displays general information about OSPFv3 routing processes. |
| timers pacing lsa-group, on page 1337 | Changes the interval at which OSPFv3 link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged |
| timers pacing retransmission , on page 1339 | Configures LSA retransmission packet pacing. |

timers pacing lsa-group

To change the interval at which Open Shortest Path First Version 3 (OSPFv3) link-state advertisements (LSAs) are collected into a group and refreshed, checksummed, or aged, use the **timers pacing lsa-group** command in an appropriate configuration mode. To restore the default value, use the **no** form of this command.

timers pacing lsa-group *seconds*

no timers pacing lsa-group

Syntax Description

| | |
|----------------|---|
| <i>seconds</i> | Interval (in seconds) at which LSAs are grouped and refreshed, checksummed, or aged. Range is 10 to 1800 seconds. |
|----------------|---|

Command Default

seconds : 240

OSPFv3 LSA group pacing is enabled by default.

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers pacing lsa-group** command to control the rate at which LSA updates occur so that high CPU or buffer utilization that can occur when an area is flooded with a very large number of LSAs can be reduced. The default settings for OSPFv3 packet pacing timers are suitable for the majority of deployments. Do not change the packet pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes the risks associated with changing the default timer values.

Cisco IOS XR software groups the periodic refresh of LSAs to improve the LSA packing density for the refreshes in large topologies. The group timer controls the interval used for group refreshment of LSAs;

however, this timer does not change the frequency that individual LSAs are refreshed (the default refresh rate is every 30 minutes).

The duration of the LSA group pacing is inversely proportional to the number of LSAs the router is handling. For example, if you have about 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure OSPFv3 group packet-pacing updates between LSA groups to occur in 60-second intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing lsa-group 60
```

Related Commands

| Command | Description |
|---|--|
| show ospfv3 , on page 1267 | Displays general information about OSPFv3 routing processes. |
| timers pacing flood , on page 1335 | Configures LSA flood packet pacing. |
| timers pacing retransmission , on page 1339 | Configures LSA retransmission packet pacing. |

timers pacing retransmission

To configure link-state advertisement (LSA) retransmission packet pacing, use the **timers pacing retransmission** command in an appropriate configuration mode. To restore the default retransmission packet pacing value, use the **no** form of this command.

timers pacing retransmission *milliseconds*

no timers pacing retransmission

Syntax Description

| | |
|---------------------|--|
| <i>milliseconds</i> | Time (in milliseconds) at which LSAs in the retransmission queue are paced. Range is 5 milliseconds to 100 milliseconds. |
|---------------------|--|

Command Default

milliseconds : 66

Command Modes

Router OSPFv3 configuration
OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **timers pacing retransmission** command to control interpacket spacing between consecutive link-state update packets in the OSPFv3 retransmission queue. This command controls the rate at which LSA updates occur. When an area is flooded with a very large number of LSAs, the LSA updates can result in high CPU or buffer utilization. Using this command reduces CPU or buffer utilization.

The default settings for OSPFv3 packet retransmission pacing timers are suitable for the majority of deployments. Do not change the packet retransmission pacing timers unless all other options to meet OSPFv3 packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, no guidelines exist for changing timer values; each OSPFv3 deployment is unique and should be considered on a case-by-case basis. The network operator assumes risks associated with changing the default packet retransmission pacing timer values.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to configure LSA flood pacing updates to occur in 55-millisecond intervals for OSPFv3 routing process 1:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers pacing retransmission 55
```

Related Commands

| Command | Description |
|---|---|
| show ospfv3, on page 1267 | Displays general information about OSPFv3 routing processes. |
| timers pacing flood, on page 1335 | Configures LSA flood packet pacing. |
| timers pacing lsa-group, on page 1337 | Changes the interval at which OSPFv3 LSAs are collected into a group and refreshed, checksummed, or aged. |

timers throttle lsa all (OSPFv3)

To set rate-limiting values for Open Shortest Path First Version 3 (OSPFv3) link-state advertisement (LSA) generation, use the **timers throttle lsa all** command in an appropriate configuration mode. To restore the default values, use the **no** form of this command.

timers throttle lsa all *start-interval hold-interval max-interval*

no timers throttle lsa all

Syntax Description

| | |
|-----------------------|--|
| <i>start-interval</i> | Minimum delay (in milliseconds) for the generation of LSAs. The first instance of LSA is always generated immediately upon a local OSPFv3 topology change. The generation of the next LSA is not before the start interval. Range is 0 to 600000 milliseconds. |
| <i>hold-interval</i> | Incremental time (in milliseconds). This value is used to calculate the subsequent rate limiting times for LSA generation. Range is 1 to 600000 milliseconds. |
| <i>max-interval</i> | Maximum wait time (in milliseconds) between generation of the same LSA. Range is 1 to 600000 milliseconds. |

Command Default

start-interval : 50 milliseconds
hold-interval : 200 milliseconds
max-interval : 5000 milliseconds

Command Modes

Router OSPFv3 configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration mode. |
| Release 4.3.0 | The default timers throttle lsa values were changed to: start-interval: 50 milliseconds and hold-interval: 200 milliseconds. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The “same LSA” is defined as an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. We recommend that you keep the *milliseconds* value of the **timers lsa arrival** command less than or equal to the *hold-interval* value of the **timers throttle lsa all** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to customize OSPFv3 LSA throttling so that the start interval is 200 milliseconds, the hold interval is 10,000 milliseconds, and the maximum interval is 45,000 milliseconds. The minimum interval between instances of receiving the same LSA is 2000 milliseconds.

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle lsa all 200 10000 45000
RP/0/RSP0/CPU0:router(config-ospfv3)# timers lsa arrival 2000
```

Related Commands

| Command | Description |
|--|---|
| show ospfv3, on page 1267 | Displays general information about OSPFv3 routing processes. |
| timers lsa arrival, on page 1333 | Sets the minimum interval at which the software accepts the same LSA from OSPFv3 neighbors. |

timers throttle spf (OSPFv3)

To turn on Open Shortest Path First Version 3 (OSPFv3) shortest path first (SPF) throttling, use the **timers throttle spf** command in an appropriate configuration mode. To turn off SPF throttling, use the **no** form of this command.

timers throttle spf *spf-start spf-hold spf-max-wait*

no timers throttle spf *spf-start spf-hold spf-max-wait*

Syntax Description

| | |
|---------------------|--|
| <i>spf-start</i> | Initial SPF schedule delay (in milliseconds). Range is 1 to 600000 milliseconds. |
| <i>spf-hold</i> | Minimum hold time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |
| <i>spf-max-wait</i> | Maximum wait time (in milliseconds) between two consecutive SPF calculations. Range is 1 to 600000 milliseconds. |

Command Default

spf-start : 50 milliseconds

spf-hold: 200 milliseconds

spf-max-wait: 5000 milliseconds

Command Modes

Router OSPFv3 configuration

OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration mode. |
| Release 4.3.0 | The default timers throttle spf values were changed to: spf-start: 50 milliseconds, spf-hold: 200 milliseconds, and spf-max-wait: 5000 milliseconds. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The first wait interval between SPF calculations is the amount of time (in milliseconds) specified by the *spf-start* argument. Each consecutive wait interval is twice the current hold level (in milliseconds) until the

wait time reaches the maximum time (in milliseconds) as specified by the *spf-max-wait* argument. Subsequent wait times remain at the maximum until the values are reset or a link-state advertisement (LSA) is received between SPF calculations.

**Tip**

Setting a low *spf-start* time and *spf-hold* time causes routing to switch to the alternate path more quickly if a failure occurs. However, it consumes more CPU processing time.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to change the start, hold, and maximum wait interval values to 5, 1000, and 90,000 milliseconds, respectively:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# timers throttle spf 5 1000 90000
```

trace (OSPFv3)

To specify the Open Shortest Path First Version 3 (OSPFv3) buffer size, use the **trace** command in router ospfv3 configuration mode. To return to the default value, use the **no** form of this command.

trace size *buffer_name size*

no trace size *buffer_name size*

Syntax Description

| | |
|--------------------|--|
| size | Deletes existing buffer and creates one with <i>N</i> entries. |
| <i>buffer_name</i> | Specifies a buffer from one of the 15 listed buffers. Refer Table 141: Buffer Types, on page 1345 table for details on the buffers. |
| <i>size</i> | Specifies allowed size for the selected buffer. Options are: 0, 256, 1024, 2048, 4096, 8192, 16384, 32768, and 65536. Select 0 to disable traces. |

Command Default

No default behavior or values

Command Modes

Router ospfv3 configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Trace buffers are used to store various traffic and processing events during the runtime. Large buffers can store more events. If the buffer becomes full, old entries are overwritten by the latest entries. In a large network, user may want to increase the trace buffer size to accommodate more events.

Table 141: Buffer Types

| Name | Description |
|-----------|-----------------------|
| adj | adjacency |
| adj_cycle | dbd/flood events/pkts |
| config | config events |

| Name | Description |
|-----------|---------------------|
| errors | errors |
| events | mda/rtrid/bfd/vrf |
| ha | startup/HA/NSF |
| hello | hello events/pkts |
| idb | interface |
| pkt | I/O packets |
| rib | rib batching |
| spf | spf/topology |
| spf_cycle | spf/topology detail |
| te | mpls-te |
| test | testing info |
| mq | message queue info |

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to set 1024 error trace entries:

```

RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 osp3
RP/0/RSP0/CPU0:router(config-ospfv3)#trace size errors ?
 0          disable trace
256        trace entries
512        trace entries
1024       trace entries
2048       trace entries
4096       trace entries
8192       trace entries
16384      trace entries
32768      trace entries
65536      trace entries
RP/0/RSP0/CPU0:router(config-ospfv3)#trace size errors 1024

```

transmit-delay (OSPFv3)

To set the estimated time required to send a link-state update packet on the interface, use the **transmit-delay** command in an appropriate configuration mode. To return to the default value, use the **no** form of this command.

transmit-delay *seconds*

no transmit-delay *seconds*

Syntax Description

| | |
|----------------|--|
| <i>seconds</i> | Time (in seconds) required to send a link-state update. Range is 1 to 65535 seconds. |
|----------------|--|

Command Default

1 second

Command Modes

Process configuration
 Area configuration
 Interface configuration
 Virtual-link configuration
 OSPFv3 VRF configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported under OSPFv3 VRF configuration submode. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Link-state advertisements (LSAs) in the update packet must have their ages incremented by the amount specified in the *seconds* argument before transmission. The value assigned should take into account the transmission and propagation delays for the interface.

If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. This setting has more significance on very low-speed links.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | ospf | read, write |

Examples

The following example shows how to configure a transmit delay of 3 seconds for GigabitEthernet interface 0/3/0/0:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 1
RP/0/RSP0/CPU0:router(config-ospfv3)# area 0
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# interface GigabitEthernet 0/3/0/0
RP/0/RSP0/CPU0:router(config-ospfv3-ar-if)# transmit-delay 3
```

Related Commands

| Command | Description |
|---|--|
| show ospfv3, on page 1267 | Displays general information about OSPF routing processes. |

virtual-link (OSPFv3)

To define an Open Shortest Path First Version 3 (OSPFv3) virtual link, use the **virtual-link** command in area configuration mode. To remove a virtual link, use the **no** form of this command.

virtual-link *router-id*

no virtual-link

Syntax Description

| | |
|------------------|---|
| <i>router-id</i> | Router ID associated with the virtual link neighbor. The router ID appears in the show ospfv3 display. This value must be entered in 32-bit dotted-decimal notation, similar to an IP Version 4 (IPv4) address. There is no default. |
|------------------|---|

Command Default

No virtual links are defined.

Command Modes

Area configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

In OSPFv3, when there exists a path through another non-backbone area over which the virtual link can function, all areas must be connected to a backbone area. If the connection to the backbone is lost, it can be repaired by establishing a virtual link.

Virtual links, which are defined in the submode of the area they transit, are in effect virtual point-to-point interfaces belonging to area 0 (the backbone). The virtual links inherit parameter values from the backbone area, rather than the transit area in which they are defined.

Each virtual link neighbor must include the router ID of the virtual link neighbor for the link to be properly established. Use the **show ospfv3** command to display the router ID of an OSPFv3 process.

Use the **virtual-link** command to place the router in virtual-link configuration mode (config-router-ar-vl), from which you can configure virtual-link-specific settings. Commands configured under this mode (such as the **transmit-delay** command) are automatically bound to that virtual link.

Task ID

| Task ID | Operations |
|---------|-------------|
| ospf | read, write |

Examples

The following example shows how to establish a virtual link with default values for all optional parameters:

```
RP/0/RSP0/CPU0:router(config)# router ospfv3 201
RP/0/RSP0/CPU0:router(config-ospfv3)# area 1
RP/0/RSP0/CPU0:router(config-ospfv3-ar)# virtual-link 10.3.4.5
```

Related Commands

| Command | Description |
|---|---|
| show ospfv3, on page 1267 | Displays general information about OSPF routing processes. |
| transmit-delay (OSPFv3), on page 1347 | Sets the estimated time required to send a link-state update packet on the interface. |

vrf (OSPFv3)

To configure an Open Shortest Path First Version 3 (OSPFv3) VPN routing and forwarding (VRF) instance and enter OSPFv3 VRF configuration submenu, use the **vrf** command in router configuration mode. To terminate an OSPFv3 VRF, use the **no** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

| | |
|-----------------|--|
| <i>vrf-name</i> | Specifies a name for the OSPFv3 vrf. If a name is not specified, the default vrf is assumed. |
|-----------------|--|

Command Default

No OSPFv3 VRF is configured.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to explicitly configure a VRF. This command creates a separate instance of OSPFv3 within the process. Commands configured under the VRF configuration mode (such as the **interface** [OSPFv3] and **authentication** [OSPFv3] commands) are automatically bound to that VRF.

To modify or remove the VRF, the *vrf-name* argument format must be the same as the format used when creating the VRF.

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to configure VRF *vrf_1* and enter OSPFv3 VRF configuration submenu:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospfv3 osp3
```

```
RP/0/RSP0/CPU0:router(config-ospfv3)#vrf vrf_1
RP/0/RSP0/CPU0:router(config-ospfv3-vrf)#
```

Related Commands

| Command | Description |
|--|---|
| router-id (OSPFv3), on page 1261 | Configures a router ID for an OSPFv3 process. |



RIB Commands

This module describes the commands used to display and clear information in the Routing Information Base (RIB).

For detailed information about RIB concepts, configuration tasks, and examples, see the *Implementing RIB on Cisco ASR 9000 Series Router* module in *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family next-hop dampening disable](#), page 1355
- [clear route](#), page 1356
- [maximum prefix \(RIB\)](#), page 1358
- [lcc](#), page 1360
- [rcc](#), page 1362
- [recursion-depth-max](#), page 1364
- [router rib](#), page 1365
- [rump always-replicate](#), page 1366
- [show lcc statistics](#), page 1368
- [show rcc](#), page 1370
- [show rcc statistics](#), page 1372
- [show rcc vrf](#), page 1374
- [show rib](#), page 1376
- [show rib afi-all](#), page 1378
- [show rib attributes](#), page 1380
- [show rib client-id](#), page 1382
- [show rib clients](#), page 1384
- [show rib extcomms](#), page 1386
- [show rib firsthop](#), page 1388
- [show rib history](#), page 1390

- [show rib next-hop](#), page 1392
- [show rib opaques](#), page 1394
- [show rib protocols](#), page 1397
- [show rib recursion-depth-max](#), page 1399
- [show rib statistics](#), page 1401
- [show rib tables](#), page 1404
- [show rib trace](#), page 1406
- [show rib vpn-attributes](#), page 1408
- [show rib vrf](#), page 1410
- [show route](#), page 1412
- [show route backup](#), page 1418
- [show route best-local](#), page 1421
- [show route connected](#), page 1423
- [show route local](#), page 1425
- [show route longer-prefixes](#), page 1427
- [show route next-hop](#), page 1430
- [show route quarantined](#), page 1433
- [show route resolving-next-hop](#), page 1436
- [show route static](#), page 1438
- [show route summary](#), page 1440

address-family next-hop dampening disable

To disable Routing Information Base (RIB) next-hop dampening, use the **address-family next-hop dampening disable** command in router configuration mode. To enable RIB next-hop dampening, use the **no** form of this command.

address-family {ipv4| ipv6} next-hop dampening disable

no address-family {ipv4| ipv6} next-hop dampening disable

Syntax Description

| | |
|------|---|
| ipv4 | Specifies IP Version 4 (IPv4) address prefixes. |
| ipv6 | Specifies IP Version 6 (IPv6) address prefixes. |

Command Default

RIB next-hop dampening is enabled.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| rib | read, write |

Examples

The following example shows how to disable RIB next-hop dampening for IPv6 address families:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router rib
RP/0/RSP0/CPU0:router(config-rib)# address-family ipv6 next-hop dampening disable
```

clear route

To clear routes from the IP routing table, use the **clear route** command in EXEC configuration mode.

```
clear route [vrf {vrf-name| all}] {ipv4| ipv6| afi-all} {unicast| multicast| safi-all} [topology topo-name]
[ip-address mask]
```

Syntax Description

| | |
|---|--|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| afi -all | Specifies IP Version 4 and IP Version 6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| safi-all | Specifies unicast and multicast address prefixes. |
| topology <i>topo-name</i> | (Optional) Specifies topology table information and name of the topology table. |
| <i>ip-address node-id</i> | (Optional) Clears hardware resource counters from the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |
| <i>ip-address</i> | Network IP address about which routing information should be displayed. |
| <i>mask</i> | Network mask specified in either of two ways: Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |

Command Default

If a **vrf** *vrf-name* is not specified, routes are cleared from the default IPv4 unicast VRF.

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear route** command to clear routes from an IP routing table to a specific network, a matching subnet address, or all routes.

Task ID

| Task ID | Operations |
|---------|-------------|
| rib | read, write |

Examples

The following example shows how to remove all routes matching the subnet address 192.168.2.0 and mask 255.255.255.0 from the IPv4 unicast routing table:

```
RP/0/RSP0/CPU0:router# clear route ipv4 unicast 192.168.2.0 255.255.255.0
```

The following example shows how to remove all routes from the IPv4 unicast routing table:

```
RP/0/RSP0/CPU0:router# clear route ipv4 unicast
```

Related Commands

| Command | Description |
|---|--|
| show route , on page 1412 | Displays the current state of the routing table. |

maximum prefix (RIB)

To set the prefix limit for the VPN routing and forwarding (VRF) instance, use the **maximum prefix** command in global VRF address family configuration mode. To set the prefix limits to the default values, use the **no** form of this command.

maximum prefix *maximum* [*mid-threshold*]

no maximum prefix

Syntax Description

| | |
|----------------------|--|
| <i>maximum</i> | Maximum number of prefixes allowed in the VRF instance. Range is 32 to 2000000. |
| <i>mid-threshold</i> | (Optional) Integer specifying at what percentage of the <i>maximum</i> argument value the software starts to generate a Simple Network Management Protocol (SNMP) trap. Range is 1 to 100. |

Command Default

No default behavior or values

Command Modes

Global VRF address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **maximum prefix** command to configure a maximum number of prefixes that a VRF instance is allowed to receive.

A maximum number of routes is applicable to dynamic routing protocols as well as static or connected routes.

Task ID

| Task ID | Operations |
|---------|-------------|
| rib | read, write |

Examples

The following example shows how to set the maximum number of prefixes allowed to 1000:

```
RP/0/RSP0/CPU0:router(config)# vrf vrf-A
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# maximum prefix 1000
```

Related Commands

| Command | Description |
|---|---------------------------------------|
| show rib tables, on page 1404 | Displays all tables known to the RIB. |

lcc

To enable Label Consistency Checker (lcc) background scan for IPv6 or IPv4 labels, use the **lcc enable** command in global configuration mode. To disable lcc background scan, use the **no** for of this command.

lcc {ipv4| ipv6} unicast {enable| period *milliseconds*}

no lcc {ipv4| ipv6} unicast {enable| period *milliseconds*}

Syntax Description

| | |
|-----------------------------------|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| period <i>milliseconds</i> | Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds. |

Command Default

Label consistency checker is disabled.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| ipv4 | read, write |
| ipv6 | read, write |

Examples

This example shows how to enable lcc for IPv6 labels:

```
RP/0/RSP0/CPU0:router#configure
```

```
RP/0/RSP0/CPU0:router(config)#lcc ipv6 unicast enable
```

rcc

To enable Route Consistency Checker (rcc) background scan for IPv6 or IPv4 routes, use the **rcc enable** command in global configuration mode. To disable rcc background scan, use the **no** form of this command.

```
rcc {ipv4| ipv6} unicast {enable| period milliseconds}
```

```
no rcc {ipv4| ipv6} unicast {enable| period milliseconds}
```

Syntax Description

| | |
|-----------------------------------|--|
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| period <i>milliseconds</i> | Specifies the period between scans in milliseconds. Range is 100 to 600000 milliseconds. |

Command Default

Route consistency checker is disabled.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **period** option to control how often the scan be triggered. Each time the scan is triggered, the background scan process resumes verification from where it was left out and sends one buffer's worth of routes to the forwarding information base (FIB).

Task ID

| Task ID | Operation |
|---------|-------------|
| ipv4 | read, write |
| ipv6 | read, write |

Examples

This example shows how to configure rcc for IPv6 unicast:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#rcc ipv6 unicast enable
```

This example shows how to enable rcc with a scan period of 500 milliseconds for IPv6 unicast:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#rcc ipv6 unicast period 500
```

recursion-depth-max

To set the maximum depth for route recursion checks, use the **recursion-depth-max** command in router configuration mode. To set the recursion checks to the default value, use the **no** form of this command.

recursion-depth-max *maximum*

no recursion-depth-max *maximum*

| Syntax Description | <i>maximum</i> | Maximum depth for recursion checks. Range is 5 to 16. |
|--------------------|----------------|---|
|--------------------|----------------|---|

Command Default The default recursion depth is 128.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **recursion-depth-max** command to configure a specific maximum number of recursion checks in the range of 5 to 16.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rib | read, write |

Examples The following example shows how to set the maximum depth for route recursion checks to 12:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router rib
RP/0/RSP0/CPU0:router(config-rib)# recursion-depth-max 12
```

router rib

To enter Routing Information Base (RIB) configuration mode, use the **router rib** command in global configuration mode. To remove all RIB configurations and terminate the RIB routing process, use the **no** form of this command.

router rib

no router rib

Syntax Description This command has no arguments or keywords.

Command Default Router configuration mode is not enabled.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | bgp | read, write |
| | ospf | read, write |
| | hsrp | read, write |
| | isis | read, write |

Examples The following example shows how to enter RIB configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router rib
```

rump always-replicate

To enable replication from uRIB to muRIB as usual even after features such as MTR are configured, use the **rump always-replicate** command in router configuration mode. To disable replication from uRIB to muRIB, use the **no** form of this command.

rump always-replicate [*access-list*]

no rump always-replicate [*access-list*]

Syntax Description

| | |
|-------------------------|-------------------------------------|
| <i>access-list-name</i> | (Optional) Name of the access list. |
|-------------------------|-------------------------------------|

Command Default

Replication from uRIB to muRIB is enabled.

Command Modes

Router address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.9.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Configuring the **rump always-replicate** command allows routers in a network to be upgraded to multitopology routing gradually without a flag day where all routers need to be configured at the same time without major service disruption. When **rump always-replicate** is configured, replicated routes are added into the muRIB with the lowest admin distance. So if protocols are populating the muRIB, they continue to do so. For the same route, protocol routes win over replicated routes because of higher admin distance.

If an unwanted more specific route comes from the uRIB, optionally provide an access list through which the replicated routes are run. If the route passes the access list, the route is replicated by RUMP.

Task ID

| Task ID | Operations |
|---------|-------------|
| rib | read, write |

Examples

The following example shows how to enable replication from uRIB to muRIB:

```
RP/0/RSP0/CPU0:router(config)# rump rib
```

```
RP/0/RSP0/CPU0:router(config-rib)# address-family ipv4  
RP/0/RSP0/CPU0:router(config-rib-afi)# rump always-replicate
```

show lcc statistics

To view results of a label consistency checker (lcc) background scan, use the **show lcc statistics** command in EXEC mode.

show lcc {ipv4|ipv6} unicast statistics

Syntax Description

| | |
|----------------|-----------------------------------|
| ipv4 | IPv4 address prefix. |
| ipv6 | IPv6 address prefix. |
| unicast | Specifies unicast address prefix. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ipv4 | read |
| ipv6 | read |

Examples

This example shows background scan statistics for AFI-SAFI mpls6-unicast:

```
RP/0/RSP0/CPU0:router#show lcc ipv6 unicast statistics
```

```
Background Scan Statistics for AFI-SAFI mpls6-unicast:
```

```
=====
```

```
Scan enabled:           False
Current scan-id:       0
Scan triggered:        False
```

```

Configured period:      60          Current period:      0
Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries          Damping percent:    70
Default route churn:  10          Current route churn: 0
Route churn last calculated at    Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI mpls6-unicast:
=====

```

End Of Logs

This example shows background scan statistics for AFI-SAFI mpls4-unicast:

```
RP/0/RSP0/CPU0:router#show lcc ipv4 unicast statistics
```

```
Background Scan Statistics for AFI-SAFI mpls4-unicast:
=====
```

```

Scan enabled:           False
Current scan-id:       0          Scan triggered:      False
Configured period:    60          Current period:      0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries          Damping percent:    70
Default route churn:  10          Current route churn: 0
Route churn last calculated at    Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI mpls4-unicast:
=====

```

End Of Logs

show rcc

To display route consistency checker (RCC) information, use the **show rcc** command in EXEC mode.

```
show rcc {ipv4| ipv6} unicast [log| prefix netmask vrf vrf-name]
```

Syntax Description

| | |
|----------------------------|---|
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| log | (Optional) Specifies the RCC log. |
| prefix | (Optional) Starting prefix. |
| netmask | (Optional) Network mask. |
| vrf <i>vrf-name</i> | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ipv4 | read |

Examples

The following is sample output from the **show rcc** command:

```
RP/0/RSP0/CPU0:router# show rcc ipv4 unicast log
ipv4-unicast: disabled, count = 1000, period = 60, table wraps = 0
-----
node                checks performed          errors
0/6/CPU0            0                          0
0/4/CPU1            30                         0
0/4/CPU0            0                          0
0/1/CPU0            0                          0
0/RP1/CPU0          120                        0
0/RSP0/CPU0        z 0                        0
```

show rcc statistics

To view results of a route consistency checker (rcc) background scan, use the **show rcc statistics** command in EXEC mode.

show rcc {ipv4| ipv6} unicast statistics

Syntax Description

| | |
|----------------|-------------------------------------|
| ipv4 | IPv4 address prefix. |
| ipv6 | IPv6 address prefix. |
| unicast | Specifies unicast address prefixes. |

Command Default

None

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ipv4 | read |
| ipv6 | read |

Examples

This example shows background scan statistics for AFI-SAFI IPv6 unicast:

```
RP/0/RSP0/CPU0:router#show rcc ipv6 unicast statistics
```

```
Background Scan Statistics for AFI-SAFI ipv6-unicast:
```

```
=====
```

```
Scan enabled:           False
Current scan-id:       0
Scan triggered:        False
```

```
Configured period:      60           Current period:      0
Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries           Damping percent:    70
Default route churn:  10           Current route churn: 0
Route churn last calculated at      Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI ipv6-unicast:
=====
```

End Of Logs

This example shows background scan statistics for AFI-SAFI Ipv4 unicast:

```
RP/0/RSP0/CPU0:router#show rcc ipv4 unicast statistics
```

```
Background Scan Statistics for AFI-SAFI ipv4-unicast:
=====
```

```
Scan enabled:           False
Current scan-id:        0           Scan triggered:      False
Configured period:     60           Current period:      0

Paused by range scan: False
Paused by route churn: False
Paused by error scan: False

Last data sent: 0 entries           Damping percent:    70
Default route churn:  10           Current route churn: 0
Route churn last calculated at      Dec 31 16:00:00.000

Logs stored for background scan ids:

Log for AFI-SAFI ipv4-unicast:
=====
```

End Of Logs

show rcc vrf

To run on-demand route consistency checker (rcc) scan on AFI, SAFI, table, and prefix or the entire set of prefixes in the table, use the **show rcc vrf** command in EXEC mode.

```
show rcc {ipv4|ipv6} unicast prefix/mask vrf vrfname
```

Syntax Description

| | |
|---------------------|--|
| ipv4 | IPv4 address prefix. |
| ipv6 | IPv6 address prefix. |
| <i>prefix /mask</i> | Specifies unicast address prefix. |
| vrf | Specifies VPN routing and forwarding (VRF) instance. |
| <i>vrfname</i> | Name of the VRF. |

Command Default

None.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| ipv4 | read |
| ipv6 | read |

Examples

This example shows how to run on-demand rcc scan for an IPv6 prefix:

```
RP/0/RSP0/CPU0:router#show rcc ipv6 unicast 2001:DB8::/32 vrf vrf_1
```

This example shows how to run on-demand rcc scan for an Ipv4 prefix:

```
RP/0/RSP0/CPU0:router#show rcc ipv4 unicast 10.2.3.4/32 vrf vrf-1
```

show rib

To display Routing Information Base (RIB) data, use the **show rib** command in EXEC mode.

```
show rib {ipv4|ipv6} {unicast|multicast}[firsthop [ type interface-path-id]| next-hop [ type interface-path-id]| opaques | {attribute | ip-nexthop| ipfrr| safi-tunnel| summary | tunnel-nexthop}]| protocols [ standby]| statistics [ name]| [standby]| topology | {topo-name| all}]
```

Syntax Description

| | |
|-------------------------------|---|
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| firsthop | (Optional) Specifies registered first-hop notification addresses. |
| type | Interface type. For more information, use the question mark (?) online help function. |
| interface-path-id | Identifies a physical interface or a virtual interface. Note Use the show interfaces command to see a list of all possible interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| next-hop | (Optional) Specifies registered next-hop notification addresses. |
| opaques | (Optional) Specifies opaque data installed in the RIB. |
| attribute | (Optional) Specifies opaque attributes installed in the RIB. |
| ip-nexthop | (Optional) Specifies P next-hop data installed in the RIB. |
| safi-tunnel | (Optional) Specifies subaddress family (SAFI) tunnel opaque data installed in the RIB. |
| summary | (Optional) Specifies a summary of opaque data installed in the RIB. |
| tunnel-nexthop | (Optional) Specifies tunnel next-hop opaque data installed in the RIB. |
| protocols | (Optional) Specifies registered protocols. |
| statistics <i>name</i> | (Optional) Specifies RIB statistics of a given name. |
| standby | (Optional) Specifies standby information. |
| all | (Optional) Specifies that all topology table information should be displayed. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ipv4 | read |

Examples The following example illustrates the **show rib** command:

```
RP/0/RSP0/CPU0:router# show rib
      ipv4 multicast
      topology
      BLUE
RP/0/RSP0/CPU0:router# show rib topology BLUE ipv4 multicast protocols
Protocol  Handle  Instance
isis      0         mt
```

Related Commands

| Command | Description |
|--|--|
| show rib afi-all, on page 1378 | Displays both IPv4 and IPv6 RIB information. |

show rib afi-all

To display Routing Information Base (RIB) data for both IPv4 and IPv6 address families, use the **show rib afi-all** command in EXEC mode.

show rib afi-all [attributes] [client-id] [clients] [extcomms] [firsthop] [history] [multicast] [next-hop] [opaques] [protocols] [recursion-depth-max] [safi-all] [statistics] [tables] [trace] [unicast] [vpn-attributes]

Syntax Description

| | |
|----------------------------|--|
| attributes | (Optional) Displays all BGP attributes installed in RIB. |
| client-id | (Optional) Displays RIB client ID for longer history of redistributed routes sent to the client. |
| clients | (Optional) Displays RIB clients. |
| extcomms | (Optional) Displays all extended communities installed in RIB. |
| firsthop | (Optional) Displays registered firsthop notification addresses. |
| history | (Optional) Displays redistributed routes sent to RIB clients. |
| multicast | (Optional) Displays multicast commands. |
| next-hop | (Optional) Displays registered next-hop notification addresses. |
| opaques | (Optional) Displays opaque data installed in RIB. |
| protocols | (Optional) Displays registered protocols. |
| recursion-depth-max | (Optional) Displays maximum recursion depth in RIB. |
| safi-all | (Optional) Displays unicast and multicast commands. |
| statistics | (Optional) Displays RIB statistics. |
| tables | (Optional) Displays a list of tables known to RIB. |
| trace | (Optional) Displays RIB trace entries. |
| unicast | (Optional) Displays unicast commands. |
| vpn-attributes | (Optional) Displays all VPN attributes installed in RIB. |

Command Default

No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | ipv4 | read |

Examples The following example illustrates the **show rib afi-all attributes** command:

```
RP/0/RSP0/CPU0:router# show rib afi-all attributes
BGP attribute data in IPv4 RIB:
0 Attributes, for a total of 0 bytes.
BGP attribute data in IPv6 RIB:
0 Attributes, for a total of 0 bytes.
```

Related Commands

| Command | Description |
|--|---------------------------|
| show rib, on page 1376 | Displays RIB information. |

show rib attributes

To display Border Gateway Protocol (BGP) attributes installed in the Routing Information Base (RIB), use the **show rib attributes** command in EXEC mode.

show rib attributes [summary] [standby]

| Syntax Description | |
|--------------------|---|
| summary | (Optional) Displays a summary of BGP attribute data installed in the RIB. |
| standby | (Optional) Displays standby information. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples The following is sample output from the **show rib attributes** command:

```
RP/0/RSP0/CPU0:router# show rib attributes
BGP attribute data in IPv4 RIB:
Attribute ID (0x2):size (68)
Attribute ID (0x3):size (52)
Attribute ID (0x4):size (68)
Attribute ID (0x5):size (52)

4 Attributes, for a total of 240 bytes.

Attribute ID : ID assigned for the attribute by BGP
```

size : size of the attribute data.

show rib client-id

To display Routing Information Base (RIB) redistribution histories, use the **show rib client-id** command in EXEC mode.

show rib client-id *id* redistribution history [standby]

Syntax Description

| | |
|------------------------|--|
| <i>id</i> | ID of the client. Range is 0 to 4294967295. |
| redistribution history | Displays longer history of redistributed routes sent to RIB clients. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib client-id** command to display a history of the route additions, deletions, and updates sent from RIB to the client across VRFs.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib client-id** command:

```
RP/0/RSP0/CPU0:router# show rib client-id 13 redistribution history
PID      JID    Client           Location
151630   113   bcd1_agent       node0_5_CPU0
  Table ID: 0xe0000000
    S 80.80.80.0/24[1/0]           update, 5 path(s), 0x0   Jan 31 09:54:57.224
```

```

S 80.80.80.0/24[1/0]          update, 6 path(s), 0x0   Jan 31 09:53:39.736
S 140.140.140.0/24[1/0]     update, 1 path(s), 0x0   Jan 31 09:53:39.729
S 80.80.80.0/24[1/0]       update, 5 path(s), 0x0   Jan 30 22:08:38.551
S 140.140.140.0/24         deleted,                  Jan 30 22:08:38.543
S 80.80.80.0/24[1/0]       update, 6 path(s), 0x0   Jan 30 22:03:05.889
S 100.100.100.0/24[1/0]    update, 1 path(s), 0x0   Jan 30 22:03:05.880

```

This table describes the significant fields shown in the display.

Table 142: show rib client-id Field Descriptions

| Field | Description |
|----------|---|
| PID | Process ID of the client. |
| JID | Job ID of the client. |
| Client | Client name. |
| Location | Location node on which the client is present. |

Related Commands

| Command | Description |
|--|-----------------------|
| show rib clients, on page 1384 | Displays RIB clients. |

show rib clients

To display Routing Information Base (RIB) clients, use the **show rib clients** command in EXEC mode.

```
show rib [afi-all| ipv4| ipv6] clients [protocols| redistribution [history]] [standby]
```

Syntax Description

| | |
|-----------------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| protocols | (Optional) Specifies client protocols. |
| redistribution | (Optional) Specifies protocols redistributed by clients |
| history | (Optional) Specifies redistributed routes sent to RIB clients. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib clients** command to display the list of clients who have registered with RIB, what protocol routes they are redistributing, and a history of the routes sent to the client.

The maximum number of redistribution entries is 5000 for Bulk Content Downloader (BCDL) and 500 for other protocols.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib clients** command:

```
RP/0/RSP0/CPU0:router# show rib clients

Process          Location          Client ID  Redist  Proto
isis             node0_5_CPU0     0          insync  insync
ospf             node0_5_CPU0     2          insync  insync

RP/0/RSP0/CPU0:router# show rib clients redistribution

isis node0_5_CPU0
  ipv4 uni         vrf default     insync          route
  static                               insync
ospf node0_5_CPU0
  ipv4 uni         vrf default     insync          route
  static                               insync
  local                               insync
bgp node0_5_CPU0
  ipv4 uni         vrf abc         insync          route
  static                               insync
bcdl_agent node0_5_CPU0
  ipv4 uni         vrf default     insync          rib_fib
  ipv4 uni         vrf bar         insync          rib_fib
  ipv4 uni         vrf abc         insync          rib_fib
  ipv4 uni         vrf test        insync          rib_fib
```

This table describes the significant fields shown in the display.

Table 143: show rib clients Field Descriptions

| Field | Description |
|-----------|--|
| Process | Client process name. |
| Location | Location where the client process is running. |
| Client ID | ID assigned to the client by RIB. |
| Redist | Whether the client is redistributing any protocols or not and whether it has read all routes from RIB or not. <ul style="list-style-type: none"> • insync—read • outsync—not read. |
| Proto | Whether the protocol has sent all its routes to RIB and signaled update complete or not. <ul style="list-style-type: none"> • insync—read • outsync—not read. |

show rib extcomms

To display all extended communities installed in the Routing Information Base (RIB), use the **show rib extcomms** command in EXEC mode.

show rib [afi-all] ipv4 ipv6 extcomms [summary] [standby]

Syntax Description

| | |
|----------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| summary | (Optional) Specifies a summary of all extended communities in the RIB. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib extcomms** command:

```
RP/0/RSP0/CPU0:router# show rib extcomms
Extended community data in RIB:
```

```

Extended community                Ref count
COST:128:128:41984                1
EIGRP route-info:0x8000:0        1
EIGRP AD:1:25600                  1
EIGRP RHB:255:0:16384            1
EIGRP LM:0x0:1:4470              1

```

This table describes the significant fields shown in the display.

Table 144: show rib extcomms Field Descriptions

| Field | Description |
|--------------------|---|
| Extended Community | Type of extended communities. Different protocols can add different extended communities. |
| Ref Count | Number of routes referring to the Extended community. |

show rib firsthop

To display registered first-hop notification addresses, use the **show rib firsthop** command in EXEC mode.

show rib [**vrf** {*vrf-name* | **all**}] [**afi-all** **ipv4**| **ipv6**] [**unicast**| **multicast**| **safi-all**] **firsthop** [*client-name*] [*type interface-path-id*] *ip-address* [*prefix-length*] *ip-address mask* [**resolved**| **unresolved**| **damped**] [**summary**] [**standby**]

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>ip-address</i> | (Optional) Network that BGP advertises. |
| <i>/ prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| <i>ip-address mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| resolved | (Optional) Specifies resolved next-hops. |
| unresolved | (Optional) Specifies unresolved next-hops. |
| damped | (Optional) Specifies next-hops that are damped. |

| | |
|----------------|---|
| summary | (Optional) Specifies a summary of the next-hop information. |
| standby | (Optional) Displays standby information. |

Command Default If a **vrf** *vrf-name* is not specified, the registered first-hop notifications addresses are displayed for the default IPv4 unicast VRF.

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib firsthop** command to display the list of first hops registered by various clients with RIB and the address and interface through which they are resolved.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rib | read |

Examples The following is sample output from the **show rib firsthop** command:

```
RP/0/RSP0/CPU0:router# show rib firsthop

Registered firsthop notifications:
0.0.0.0/0 via 1.1.0.1 - MgmtEth0/5/CPU0/0, ospf/node0_5_CPU0
1.1.0.1/32 via 1.1.0.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
1.1.1.1/32 via 1.1.1.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
10.10.10.1/32 via 10.10.10.1 - Loopback0, ipv4_static/node0_5_CPU0
10.10.10.3/32 via 10.10.10.3 - Loopback0, ipv4_static/node0_5_CPU0
15.15.15.1/32 via 10.10.10.1 - Loopback0, ipv4_static/node0_5_CPU0
20.20.20.1/32 via 1.1.1.1 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
30.30.30.1/32 via 1.1.1.2 - MgmtEth0/5/CPU0/0, ipv4_static/node0_5_CPU0
```

show rib history

To display history information for Routing Information Base (RIB) clients, use the **show rib history** command in EXEC mode.

show rib [afi-all] ipv4|ipv6] history [client-id *client-id*] [standby]

Syntax Description

| | |
|-----------------------------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| client-id <i>client-id</i> | (Optional) Specifies the ID of the client. Range for <i>client-id</i> argument is 0 to 4294967295. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib history** command to display the list of routes that RIB has sent to various clients.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib history** command:

```
RP/0/RSP0/CPU0:router# show rib history
```

```

JID  Client          Location
229  isis             node0_5_CPU0
    Table ID: 0xe0000000
    S 80.80.80.0/24[1/0]    update, 6 path(s),    04:32:09
    S 100.100.100.0/24[1/0] update, 1 path(s),    04:32:09
    S 40.40.40.0/24[1/0]   update, 1 path(s),    04:32:09
    S 15.15.15.0/24[1/0]   update, 1 path(s),    04:32:09
JID  Client          Location
260  ospf             node0_5_CPU0
    Table ID: 0xe0000000
    S 80.80.80.0/24[1/0]    update, 6 path(s),    04:32:09
    S 100.100.100.0/24[1/0] update, 1 path(s),    04:32:09
    S 40.40.40.0/24[1/0]   update, 1 path(s),    04:32:09
    S 15.15.15.0/24[1/0]   update, 1 path(s),    04:32:09

```

This table describes the significant fields shown in the display.

Table 145: show rib history Field Descriptions

| Field | Description |
|----------|--|
| JID | Job ID of the client process. |
| Client | Name of the client process. |
| Location | Information about where the client process is running. |

show rib next-hop

To display registered next-hop notification addresses, use the **show rib next-hop** command in EXEC mode.

```
show rib [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| safi-all] next-hop [client-name ]
[type interface-path-id] ip-address [prefix-length] ip-address mask resolved| unresolved| damped] [summary]
[standby]
```

Syntax Description

| | |
|---|--|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |

| | |
|------------------------|---|
| <i>/ prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| resolved | (Optional) Specifies resolved next-hops. |
| unresolved | (Optional) Specifies unresolved next-hops. |
| damped | (Optional) Specifies next-hops that are damped. |
| summary | (Optional) Specifies a summary of the next-hop information. |
| standby | (Optional) Displays standby information. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib next-hop** command to display the list of next-hops registered by various clients with the RIB and the address and interface through which they are resolved.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rib | read |

Examples The following is sample output from the **show rib next-hop** command:

```
RP/0/RSP0/CPU0:router# show rib next-hop
Registered nexthop notifications:
0.0.0.0/0 via 172.29.52.1 - MgmtEth0/RP1/CPU0/0, ospf/node0_RP0_CPU0
172.29.52.1/32 via 172.29.52.1 - MgmtEth0/RP1/CPU0/0, ipv4_static/node0_RP0_CPU0
```

show rib opaques

To display opaque data installed in the Routing Information Base (RIB), use the **show rib opaques** command in EXEC mode.

```
show rib [vrf {vrf-name| all}] [afi-all ipv4| ipv6] [unicast| multicast| safi-all] opaques {attribute|
ip-nexthop| ipfrr| safi-tunnel| summary| tunnel-nexthop} [ rib-client-name ] [standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| attribute | Displays opaque attributes installed in the RIB. |
| ip-nexthop | Displays IP next-hop data installed in the RIB. |
| ipfrr | Displays IP fast reroute (IPFRR) opaque data installed in the RIB. |
| safi-tunnel | Displays subaddress family (SAFI) tunnel opaque data installed in the RIB. |
| summary | Displays a summary of opaque data installed in the RIB. |
| tunnel-nexthop | Displays tunnel next-hop opaque data installed in the RIB. |
| <i>rib-client-name</i> | (Optional) Name of the RIB client. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If information is not used by the RIB server process, it is viewed as opaque data. Use the **show rib opaques** command to display opaque data installed in the RIB.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib opaques** command:

```
RP/0/RSP0/CPU0:router# show rib opaques safi-tunnel

Summary of safi tunnel opaque data in IPv4 RIB:

Opaque key: 1:10.1.0.2
Opaque data:
Tunnel Encap - ifhandle=0x1000180, type=L2TPv3, Params=[Session-id=0x1EB1127C, `
CookieIen=8, Cookie=0xA73A3E0AFCD419A6] Opaque key: 65535:10.0.101.1 Opaque data:

RP/0/RSP0/CPU0:router# show rib ipv6 opaques tunnel-nexthop

Summary of 6PE/6VPE IP over tunnel nexthop opaque data in IPv6 RIB:

Opaque key: 1::ffff:10.1.0.2
Opaque key: 65535::ffff:10.0.101.1
Opaque key: 65535::ffff:10.0.101.2
Opaque key: 65535::ffff:10.0.101.3
Opaque key: 65535::ffff:10.0.101.4
Opaque key: 65535::ffff:10.0.101.5
```

This table describes the significant fields shown in the display.

Table 146: show rib opaques Field Descriptions

| Field | Description |
|-------------|---|
| Opaque key | Unique key for the opaque data as populated by the protocol client. |
| Opaque data | Data for the given key. |

Related Commands

| Command | Description |
|--|--|
| show route, on page 1412 | Displays current routes information in the Routing Information Base (RIB). |
| show ospf routes, on page 1131 | Displays Open Shortest Path First (OSPF) topology table. |

show rib protocols

To display protocols registered for route addition, use the **show rib protocols** command in EXEC mode.

```
show rib [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| safi-all] protocols [standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the registered first-hop notification addresses are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib protocols** command:

```
RP/0/RSP0/CPU0:router# show rib protocols
```

```
Protocol  Handle  Instance
isis      0        rib
connected 1
static    2
local     3
bgp       4        102
ospf      5        1
```

This table describes the significant fields shown in the display.

Table 147: show rib protocols Field Descriptions

| Field | Description |
|----------|---|
| Protocol | Name of the protocol. |
| Handle | Handle assigned to the protocol instance. |
| Instance | Protocol instance. |

show rib recursion-depth-max

To display the maximum recursion depth in the Routing Information Base (RIB), use the **show rib recursion-depth-max** command in EXEC mode.

```
show rib [afi-all| ipv4| ipv6] recursion-depth-max [standby]
```

Syntax Description

| | |
|----------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib recursion-depth-max** command to display the maximum recursion depth for RIB. Recursion depth is the number of next-hops that can be specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib recursion-depth-max** command:

```
RP/0/RSP0/CPU0:router# show rib recursion-depth-max
IPv4:
-----
```

```
Maximum recursion depth in RIB:
```

```
  Configured: 12
    In Use: 128
```

```
IPv6:
```

```
-----
Maximum recursion depth in RIB:
```

```
  Configured: 12
    In Use: 128
```

This table describes the significant fields shown in the display.

Table 148: show rib recursion-depth-max Field Descriptions

| Field | Description |
|------------|--|
| Configured | Value of maximum recursion depth currently configured. |
| In Use | Value of maximum recursion depth RIB is using. This value can be different from the configured value because RIB has to be restarted after the configuration is changed for the new configuration to be effective. |

show rib statistics

To display Routing Information Base (RIB) statistics, use the **show rib statistics** command in EXEC mode.

```
show rib [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| safi-all] statistics [ client-name ]
[standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>client-name</i> | (Optional) Name of the RIB client. |
| standby | (Optional) Displays standby information. |

Command Default

If **vrf** *vrf-name* is not specified, the registered first-hop notification addresses are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib statistics** command to display RIB statistics. The statistics include requests sent from the clients to the RIB and the information redistributed to the client.

RIB maintains counters for all requests sent from a client including

- Route operations
- Table registrations
- Next-hop registrations
- Redistribution registrations
- Attribute registrations
- Synchronization completion

RIB also maintains the results of the requests.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib statistics** command:

```
RP/0/RSP0/CPU0:router# show rib statistics

RIB Statistics:
Received 142 batch messages
    137 route operations, 0 attribute operations
    0 opaque operations
    11 complete operations, 0 convergent operations
Results of the batch message received:
142 successes
0 forward references, 0 invalid client id, 0 unknown errors
0 memory allocation errors, 0 client lookup errors, table lookup errors 0
0 proto lookup errors, 0 client proto lookup errors
ipv4_connected/node0_RP0_CPU0 last performed route operation
with status BATCH_SUCCESS at Jun 26 21:43:33.601

Received 217422 light weight messages
4 route add requests, 2 route delete requests
10 protocol registered, 1 protocol unregistered
0 protocol modify, 0 protocol purged
14 protocol redistributions, 0 unregistered protocol redistributions
0 reset protocol redistributions
3 first hop registered, 1 first hop unregistered
3 advertisements, 0 unregistered advertisement
57 bind data, 97 update completes, 217230 other requests
udp/node0_RP0_CPU0 last performed firsthop lookup operation
with status success at Jun 27 10:09:59.990

Received 0 nexthop batch messages
0 successes
0 inits
0 registers, 0 unregisters
0 register complete, 0 sync unregistered, 0 batch finished
```

This table describes the significant fields shown in the display.

Table 149: show rib statistics Field Descriptions

| Field | Description |
|--|---|
| Received | Statistics received including batch messages and route, attribute, complete, and convergent operations. |
| Results of the batch message received | Batch message results. |
| Received <i>n</i> light weight messages | Number of lightweight API messages sent from RIB clients. |
| Received <i>n</i> nexthop batch messages | Number of batch API messages sent from RIB clients received by the RIB. |

show rib tables

To display all tables known to the Routing Information Base (RIB), use the **show rib tables** command in EXEC mode.

show rib [afi-all| ipv4| ipv6] tables [summary] [standby]

Syntax Description

| | |
|----------------|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| summary | (Optional) Displays summary table information. |
| standby | (Optional) Displays standby information. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rib tables** command to display all tables known to the RIB, including table attributes. Attributes include VPN routing and forwarding (VRF) instance, address family, and maximum prefix information.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib tables** command when entered without an address:

```
RP/0/RSP0/CPU0:router# show rib tables
```

Codes: N - Prefix Limit Notified, F - Forward Referenced
 D - Table Deleted, C - Table Reached Convergence

```
VRF          SAFI  Table ID    PrfxLmt    PrfxCnt  TblVersion  N  F  D  C
default     uni   0xe0000000  2000000    72        137         N  N  N  Y
default     multi 0xe0100000  2000000    0         0          N  N  N  Y
```

This table describes the significant fields shown in the display.

Table 150: show rib tables Field Descriptions

| Field | Description |
|------------|---|
| VRF | Name of the VRF instance. |
| SAFI | Subaddress family instance. |
| Table ID | ID of the RIB table. |
| PrfxLmt | Configured prefix limit for the RIB table. |
| PrfxCnt | Number of configured prefixes in the RIB table. |
| TblVersion | Tables version number. |
| N | Message sent when prefix limit is exceeded. |
| F | Forward referenced. If Y is indicated, a table has been created by RIB because a client has registered for the table, but RIB has not heard from the router space infrastructure (RSI) about the table. RSI manages the tables. |
| D | If Y is indicated, the table has been deleted in the RSI but RIB has not cleared the information. |
| C | Table reached convergence. |

show rib trace

To display all Routing Information Base (RIB) library call tracer (ltrace) entries, use the **show rib trace** command in EXEC mode.

```
show rib [afi-all| ipv4| ipv6] trace [clear| counts| event-manager| startup| sync| timing] [unique| wrapping]
[last entries] [hexdump] [reverse] [tailif] [stats] [verbose] [file name original location node-id] location
{all| node-id}
```

Syntax Description

| | |
|--|--|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| counts clear | (Optional) Displays route clear trace entries. |
| counts | (Optional) Displays counts trace entries. |
| event-manager | (Optional) Displays RIB event manager trace entries. |
| startup | (Optional) Displays RIB startup trace entries. |
| sync | (Optional) Displays client synchronization trace entries. |
| timing | (Optional) Displays timing trace entries. |
| unique | (Optional) Displays unique entries with counts. |
| wrapping | (Optional) Displays wrapping entries. |
| last entries | (Optional) Displays a specified number of the last entries. Range is 1 to 4294967295. |
| hexdump | (Optional) Displays traces in hexadecimal format. |
| reverse | (Optional) Displays the latest traces first. |
| tailif | (Optional) Displays new traces as they are added. |
| stats | (Optional) Displays statistics. |
| verbose | (Optional) Displays internal debugging information. |
| file name original location node-id | (Optional) Displays trace entries for a specific file for the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation. |

location { **all** | *node-id* } (Optional) Displays ltrace entries for the designated node. The *node-id* argument is entered in the *rack/slot/module* notation. The **all** keyword displays ltrace entries for all nodes.

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|------------|
| | rib | read |

Examples The following is sample output from the **show rib trace** command

```
RP/0/RSP0/CPU0:router# show rib trace

1784 wrapping entries (13312 possible, 0 filtered, 1784 total)
Mar 16 14:59:27.947 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t1 Create: Management thread
Mar 16 14:59:27.959 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t2 Create: Management event
manager
Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 Initialise: RIB server
Mar 16 14:59:28.346 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 Initialise: Client collection
Mar 16 14:59:28.676 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 Initialise: DB collection
Mar 16 14:59:28.693 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 Initialise: Timer tree
Mar 16 14:59:28.694 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 RUMP: Bind to sysdb
/ipc/gl/ipv4-rib/ for protocol notification
Mar 16 14:59:29.102 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t2 Initialise: Debugging routine

Mar 16 14:59:29.128 rib/ipv4_rib/rib-io 0/RSP0/CPU0 t1 Register: read, select cb functions

Mar 16 14:59:29.137 rib/ipv4_rib/rib-startup 0/RSP0/CPU0 t1 Register: cernno DLL name
lib_rib_error.dll
.
.
.
```

show rib vpn-attributes

To display all VPN attributes installed in the Routing Information Base (RIB), use the **show rib vpn-attributes** command in EXEC mode.

show rib [afi-all] ipv4| ipv6] vpn-attributes [summary] [standby]

Syntax Description

| | |
|----------------|---|
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| summary | (Optional) Displays VPN attribute information. |
| standby | (Optional) Displays standby information. |

Command Default

The default is IPv4 address prefixes.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show rib vpn-attributes** command:

```
RP/0/RSP0/CPU0:router# show rib vpn-attributes
Extended community data in RIB:
```

```

Extended community                               Ref count
COST:128:128:41984                               2
COST:128:129:42240                               2
COST:128:129:44544                               1
COST:128:129:169984                              2
COST:128:129:307200                              1
EIGRP route-info:0x0:0                           6
EIGRP route-info:0x8000:0                         2
EIGRP AD:444:25600                                2
EIGRP AD:444:25856                                2
EIGRP AD:444:28160                                1
EIGRP AD:444:51200                                1
EIGRP AD:444:153600                               2
EIGRP RHB:255:0:16384                             2
EIGRP RHB:255:1:16384                             5
EIGRP RHB:255:1:256000                            1
EIGRP LM:0x0:1:1500                               3
EIGRP LM:0x0:1:1514                              2
EIGRP LM:0x0:1:4470                               3
EIGRP AR:0:192.168.0.13                           6
EIGRP PM:11:0                                     6

```

MVPN attribute data in RIB:

```

MVPN Attribute                               Ref count
0:0:1:f4:0:0:0:0:1:1:1:1:1                   1
0:0:2:bc:0:0:0:0:1:3:3:3:3                   10
0:0:2:bc:0:0:0:0:1:3:3:3:4                   2

```

This table describes the significant fields shown in the display.

Table 151: show rib vpn-attributes Field Descriptions

| Field | Description |
|--------------------|--|
| Extended Community | Extended community added by the protocol clients. |
| Ref Count | Number of routes referring to the same extended community. |
| MVPN Attribute | Connector attribute added by BGP to support MVPNs. |
| Ref Count | Number of routes referring to the same extended community. |

show rib vrf

To display all VRF table information in the Routing Information Base (RIB), use the **show rib vrf** command in EXEC mode.

```
show rib vrf {vrf-name| all} [ipv4] [ipv6] [afi-all] [firsthop] [next-hop] [opaques] [protocols] [statistics name]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| afi-all | (Optional) Specifies all address families. |
| firsthop | (Optional) Specifies registered first-hop notification addresses |
| next-hop | (Optional) Specifies registered next-hop notification addresses. |
| opaques | (Optional) Specifies opaque data installed in the RIB. |
| protocols | (Optional) Specifies registered protocols. |
| statistics <i>name</i> | (Optional) Specifies RIB statistics for the given name. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| ipv4 | read |

Examples

The following example shows output from the **show rib vrf all statistics** command:

```
RP/0/RSP0/CPU0:router# show rib vrf all statistics
RP/0/RSP0/CPU0:router#
```

Related Commands

| Command | Description |
|--|---------------------------|
| show rib, on page 1376 | Displays RIB information. |

show route

To display the current routes in the Routing Information Base (RIB), use the **show route** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [afi-all ipv4| ipv6] [unicast| multicast | safi-all] [protocol [instance ]|
ip-address [mask ]| ip-address/prefix-length] [standby] [detail]
```

Syntax Description

| | |
|---|--|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. This is the default. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. This is the default. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>protocol</i> | (Optional) Name of a routing protocol. If you specify a routing protocol, use one of the following keywords: <ul style="list-style-type: none"> • bgp • eigrp • isis • ospf • rip • static • local • connected |
| instance | (Optional) Number or name used to identify an instance of the specified protocol. |
| <i>ip-address</i> | (Optional) Network IP address about which routing information should be displayed. |

| | |
|-----------------------|--|
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |
| <i>/prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| standby | (Optional) Displays standby information. |
| detail | (Optional) Displays detailed information for the specified prefix. |

Command Default

If a **vrf** *vrf-name* is not specified, routes are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. The input parameters and output were modified to display 4-byte autonomous system numbers and extended communities in either asplain or asdot notations. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route** command when entered without an address:

```
RP/0/RSP0/CPU0:router# show route

Codes: C - connected, S - static, R - RIP, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G - DAGR
       A - access/subscriber, (!) - FRR Backup path

Gateway of last resort is 1.0.0.1 to network 0.0.0.0

S*  0.0.0.0/0 [1/0] via 1.0.0.1, 13:14:59
C   1.0.0.0/16 is directly connected, 13:14:59, MgmtEth0/5/CPU0/0
L   1.0.14.15/32 is directly connected, 13:14:59, MgmtEth0/5/CPU0/0
C   3.2.3.0/24 is directly connected, 00:04:39, GigabitEthernet0/3/0/0
L   3.2.3.2/32 is directly connected, 00:04:39, GigabitEthernet0/3/0/0
O E2 5.2.5.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0
O E2 6.2.6.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0
C   7.2.7.0/24 is directly connected, 00:04:20, GigabitEthernet0/3/0/7
L   7.2.7.2/32 is directly connected, 00:04:20, GigabitEthernet0/3/0/7
O E2 8.2.8.0/24 [110/20] via 3.3.3.1, 00:04:20, GigabitEthernet0/3/0/0

C   10.3.0.0/16 is directly connected, 13:14:59, GigabitEthernet0/0/0/0
L   10.3.0.2/32 is directly connected, 13:14:59, GigabitEthernet0/0/0/0
```

This table describes the significant fields shown in the display.

Table 152: show route Field Descriptions

| Field | Description |
|-------------------------|---|
| S* | Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was derived from a static (candidate default). |
| [1/0] | First number in the brackets is the administrative distance of the information source; the second number is the metric for the route. |
| 1.0.0.0/16 | Address and prefix length of the remote network. |
| MgmtEthernet 0/5/CPU0/0 | Specifies the interface through which the specified network can be reached. |
| C | Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was connected. |
| L | Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was local. |

| Field | Description |
|-------------|---|
| O | Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was on-demand routing (ODR). |
| E2 | Code indicating how the route was derived. See the code legend preceding the output. In this case, the route was OSPF external type 2. |
| 8.2.8.0/24 | Address and prefix length of the remote network connected to the static route. |
| via 3.3.3.1 | Specifies the address of the next router to the remote network. |
| 13:14:59 | Specifies the last time the route was updated. |

When you specify that you want information about a particular network, more detailed statistics are displayed. The following is sample output from the **show route** command when entered with an IP address:

```
RP/0/RSP0/CPU0:router# show route 10.0.0.0

Routing entry for 10.0.0.0/16
  Known via "connected", distance 0, metric 0 (connected)
  Installed Mar 22 22:10:20.906
  Routing Descriptor Blocks
    directly connected, via GigabitEthernet0/0/0/0
      Route metric is 0
  No advertising protos.
```

Intermediate System-to-Intermediate System (IS-IS) includes an IP address typed length value (TLV) in its link-state packet (LSP) that helps identify the node injecting the route into the network. The IS-IS node uses one of its own interface addresses in this TLV. A loopback address is preferred among interfaces configured under IS-IS. When other networking devices calculate IP routes, they can store the IP address as the originator address with each route in the routing table.

The following example shows the output from the **show route** command for a specific IP address on a router configured with IS-IS. Each path that is shown under the Routing Descriptor Blocks report displays two IP addresses. The first address (10.0.0.9) is the next-hop address; the second is the originator IP address from the advertising IS-IS router.

```
RP/0/RSP0/CPU0:router# show route 10.0.0.1

Routing entry for 10.0.0.0/8
  Known via "isis", distance 115, metric 10, type level-2
  Installed Jan 22 09:26:56.210
  Routing Descriptor Blocks:
    * 10.0.0.9, from 10.0.0.9, via GigabitEthernet2/1
      Route metric is 10
  No advertising protos.
```

This table describes the significant fields shown in the display.

Table 153: show route with IP Address Field Descriptions

| Field | Description |
|----------------------------|--|
| Routing entry for | Network address and mask. |
| Known via | Indicates how the route was derived. |
| distance | Administrative distance of the information source. |
| metric | Route value assigned by the routing protocol. |
| type | IS-IS type level. |
| Routing Descriptor Blocks: | Displays the next-hop IP address followed by the information source. |
| from ... via ... | First address is the next-hop IP address, and the other is the information source. This report is followed by the interface for this route. |
| Route metric | Best metric for this Routing Descriptor Block. |
| No advertising protos. | Indicates that no other protocols are advertising the route to their redistribution consumers. If the route is being advertised, protocols are listed in the following manner: Redist Advertisers: isis p ospf 43 |

The following example illustrates the **show route** command with the **topology** *topo-name* keyword and argument specified:

```
RP/0/RSP0/CPU0:router# show route ipv4 multicast topology green

Codes: C - connected, S - static, R - RIP, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local, G - DAGR
       A - access/subscriber, (!) - FRR Backup path

Gateway of last resort is not set

i L1 10.1.102.0/24 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
i L1 10.3.3.0/24 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
i L1 192.168.0.40/32 [115/20] via 10.1.102.41, 1w4d, GigabitEthernet0/1/0/0.1
```

This example is a sample **show route detail** command output that displays path ID and backup-path ID information:

```
RP/0/RSP0/CPU0:router#show route 10.1.1.3 detail
Routing entry for 10.1.1.16/32
  Known via "ospf 2", distance 110, metric 21, type intra area
  Installed Oct 28 16:07:05.752 for 00:01:56
  Routing Descriptor Blocks
    40.1.10.1, from 10.1.1.16, via Bundle-Ether10, Protected
      Route metric is 21
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:2          Path ref count:0
      Backup path id:33
    200.40.1.101, from 10.1.1.16, via Bundle-Ether1.1, Protected
      Route metric is 21
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:1          Path ref count:0
      Backup path id:33
    100.100.2.1, from 10.1.1.16, via TenGigE0/2/0/3.1, Backup
      Route metric is 0
      Label: None
      Tunnel ID: None
      Extended communities count: 0
      Path id:33          Path ref count:2
  Route version is 0xe (14)
  No local label
  IP Precedence: Not Set
  QoS Group ID: Not Set
  Route Priority: RIB_PRIORITY_NON_RECURSIVE_LOW (6) SVD Type RIB_SVD_TYPE_LOCAL
  No advertising protos.
```

Related Commands

| Command | Description |
|--|---|
| show interfaces | Lists interface information. |
| show route summary, on page 1440 | Displays the current contents of the routing table in summary format. |
| show rib opaques, on page 1394 | Displays opaque data installed in the Routing Information Base (RIB). |
| show ospf routes, on page 1131 | Displays Open Shortest Path First (OSPF) topology table. |

show route backup

To display backup routes from the Routing Information Base (RIB), use the **show route backup** command in EXEC mode.

```
show route [vrf {vrf-name | all}] [afi-all | ipv4 | ipv6] [unicast | multicast] {topology topo-name} | safi-all]
backup [ip-address [mask] | ip-address / | prefix-length] ][standby]
```

Syntax Description

| | |
|---|--|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | (Optional) Network IP address about which backup routing information should be displayed. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address. |
| <i>/prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the backup routes from the RIB are displayed for the default IPv4 unicast VRF.

Command Modes EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route backup** command to display information about routes that have been installed into the RIB as backup routes. This command also displays information about the currently selected active route for which there is a backup.

When the **afi-all** keyword is used, the *ip-address* and *mask* arguments are not available.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route backup** command:

```
RP/0/RSP0/CPU0:router# show route backup

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local
S      172.73.51.0/24 is directly connected, 2d20h, GigabitEthernet 4/0/0/1
      Backup O E2 [110/1] via 10.12.12.2, GigabitEthernet 3/0/0/1
```

This table describes the significant fields shown in the display.

Table 154: show route backup Field Descriptions

| Field | Description |
|----------------|--|
| S | Code indicating how the route was derived. See the legend of the codes preceding the output. |
| 172.73.51.0/24 | IP address and length of the route. |

| Field | Description |
|------------------------|---|
| 2d20h | Time (in hh:mm:ss) since the route was installed in the RIB. |
| GigabitEthernet4/0/0/1 | Outbound interface for the route. |
| Backup | Identifies the entry as a backup version of the route, typically installed by a different routing protocol. |
| O | Code indicating how the route was derived. See the code legend preceding the output. |
| E2 | Code for the type of route. This code is relevant only for OSPF and IS-IS routes. The codes for an OSPF route can be: none—intra-area route IA—interarea route E1—external type 1 E2—external type 2 N1—NSSA external type 1 N2—NSSA external type 2 The codes for an IS-IS route can be: L1—level 1 L2—level 2 ia—interarea su—summary route |
| [110/1] | Distance and metric for the route. |
| 10.12.12.2 | IP address of next-hop on the route. |
| GigabitEthernet3/0/0/1 | Outbound interface for the OSPF version of the route. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current routes in the RIB. |

show route best-local

To display the best local address to use for return packets from the given destination, use the **show route best-local** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [ipv4| ipv6] [unicast| multicast| {topology topo-name}| safi-all] best-local
ip-address [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | IP address about which best local information should be displayed. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the best local address is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route best-local** command to display information about the best local routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route best-local** command:

```
RP/0/RSP0/CPU0:router# show route best-local 10.12.12.1/32
Routing entry for 10.12.12.1/32
  Known via "local", distance 0, metric 0 (connected)
  Routing Descriptor Blocks
    10.12.12.1 directly connected, via GigabitEthernet3/0/0/1
    Route metric is 0
```

This table describes the significant fields shown in the display.

Table 155: show route best-local Field Descriptions

| Field | Description |
|---|--|
| Routing entry for | Identifies the requested IP address. |
| Known via | Indicates how the route was derived. |
| distance | Administrative distance of the information source. |
| metric | Route value assigned by the routing protocol. |
| Routing Descriptor Blocks: | Displays the next-hop IP address followed by the information source. |
| 10.12.12.1 Directly connected ... via ... | First address is the next-hop IP address, followed by a report that it is directly connected. This report is followed by the interface for this route. |

Related Commands

| Command | Description |
|--|---|
| show route local, on page 1425 | Displays local addresses installed in the RIB as a receive entry. |

show route connected

To display the current connected routes of the routing table, use the **show route connected** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| {topology topo-name}] safi-all
connected [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the current connected routes of the routing table are displayed for the default IPv4 unicast VRF.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route connected** command to display information about connected routes in the routing table.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route connected** command:

```
RP/0/RSP0/CPU0:router# show route connected
C    1.68.0.0/16 is directly connected, 13:43:40, MgmtEth0/5/CPU0/0
C    3.3.3.0/24 is directly connected, 00:23:23, GigabitEthernet0/3/0/0
C    7.7.7.0/24 is directly connected, 00:33:00, GigabitEthernet0/3/0/7
C    10.0.0.0/16 is directly connected, 13:43:40, GigabitEthernet0/0/0/0
C    10.10.10.0/30 is directly connected, 13:43:40, Loopback0
C    11.11.11.0/24 is directly connected, 13:43:40, Loopback11
```

This table describes the significant fields shown in the display.

Table 156: show route connected Field Descriptions

| Field | Description |
|-------------------|--|
| C | Code to indicate the route is connected. |
| 1.68.0.0/16 | IP address and length of the route. |
| 13:43:40 | Time (in hh:mm:ss) since the route was installed in the RIB. |
| MgmtEth0/5/CPU0/0 | Outbound interface for the route. |

Related Commands

| Command | Description |
|--|---|
| show route summary, on page 1440 | Displays the current contents of the RIB. |

show route local

To display local routes receiving routing updates from the Routing Information Base (RIB), use the **show route local** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [afi-all ipv4| ipv6] [unicast| multicast| {topology topo-name}] safi-all
local [type interface -path-id] [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the local routes receiving updates from the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route local** command to display information about local routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route local** command:

```
RP/0/RSP0/CPU0:router# show route local
L   10.10.10.1/32 is directly connected, 00:14:36, Loopback0
L   10.91.36.98/32 is directly connected, 00:14:32, GigabitEthernet6/0/0/1
L   172.22.12.1/32 is directly connected, 00:13:35, GigabitEthernet3/0/0/1
L   192.168.20.2/32 is directly connected, 00:13:27, GigabitEthernet4/0/0/1
L   10.254.254.1/32 is directly connected, 00:13:26, GigabitEthernet5/0/0/1
```

This table describes the significant fields shown in the display.

Table 157: show route local Field Descriptions

| Field | Description |
|---------------|--|
| L | Code to indicate the route is local. |
| 10.10.10.1/32 | IP address and length of the route. |
| 00:14:36 | Time (in hh:mm:ss) since the route was installed in the RIB. |
| Loopback0 | Outbound interface for the route. |

Related Commands

| Command | Description |
|---|--|
| show route connected , on page 1423 | Displays information about all clients that have registered with the RIB as protocols. |

show route longer-prefixes

To display the current routes in the Routing Information Base (RIB) that share a given number of bits with a given network, use the **show route longer-prefixes** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [ipv4| ipv6] [unicast| multicast| {topology topo-name}] safi-all]
longer-prefixes {ip-address mask | ip-address/prefix-length} [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | Network IP address about which routing information should be displayed. |
| <i>mask</i> | Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |
| <i>/ prefix-length</i> | Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the current routes in the RIB sharing a specified number of bits with a network are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route longer-prefixes** command to troubleshoot forwarding problems whose cause may be a long prefix.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route longer-prefixes** command:

```
RP/0/RSP0/CPU0:router# show route longer-prefixes 172.16.0.0/8

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

L   172.29.52.70/32 is directly connected, 4d15h, MgmtEth0/RSP0/CPU0/0
L   172.29.52.71/32 is directly connected, 4d15h, MgmtEth0/RP1/CPU0/0
L   172.29.52.72/32 [0/0] via 172.29.52.72, 4d15h, MgmtEth0/RSP0/CPU0/0
```

This table describes the significant fields shown in the display.

Table 158: show route longer-prefixes Field Descriptions

| Field | Description |
|-----------------------|--|
| 172.29.52.70/32 | IP address and length of the route. |
| 4d15h | Time (in hh:mm:ss or <i>ndnh</i>) since the route was installed in the RIB. |
| MgmtEth0/RSP0 /CPU0/0 | Outbound interface for the route. |

Related Commands

| Command | Description |
|--|---|
| router static | Establishes a static route. |
| show interfaces | Lists interface information. |
| show route summary, on page 1440 | Displays the current contents of the routing table in summary format. |

show route next-hop

To filter routes by the next-hop address or interface, use the **show route next-hop** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [ipv4| ipv6] [unicast| multicast| {topology topo-name}| safi-all] next-hop
[ip-address][[standby]]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | (Optional) IP address about which next-hop information is to be displayed. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route next-hop** command to find all routes going through a next-hop address or interface.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route next-hop** command filtering routes on the next-hop address:

```
RP/0/RSP0/CPU0:router# show route next-hop 1.68.0.1

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local

Gateway of last resort is 1.68.0.1 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 1.68.0.1, 15:01:49
S 223.255.254.254/32 [1/0] via 1.68.0.1, 15:01:49
```

The following is sample output from the **show route next-hop** command filtering routes on the next-hop interface:

```
RP/0/RSP0/CPU0:router# show route next-hop GigabitEthernet 0/1/0/2

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local

Gateway of last resort is 1.68.0.1 to network 0.0.0.0

C 11.1.1.0/24 is directly connected, 15:01:46, GigabitEthernet0/1/0/2
L 11.1.1.2/32 is directly connected, 15:01:46, GigabitEthernet0/1/0/2
```

This table describes the significant fields shown in the display.

Table 159: show route next-hop Field Descriptions

| Field | Description |
|------------------------|--|
| 11.1.1.0/24 | IP address and length of the route. |
| 15:01:46 | Time (in hh:mm:ss or <i>ndnh</i>) since the route was installed in the RIB. |
| GigabitEthernet0/1/0/2 | Outbound interface for the route. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current contents of the routing table. |

show route quarantined

To display mutually recursive (looping) routes, use the **show route quarantined** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [ipv4|ipv6] [unicast|multicast] {topology topo-name}| safi-all] quarantined
[ip-address/prefix-length] ip-address mask] [standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | (Optional) IP address about which looping routes information is to be displayed. |
| <i>/ prefix-length</i> | (Optional) Length of the IP address prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash (/) must precede the decimal value. |
| <i>ip-address mask</i> | (Optional) Network mask applied to the <i>ip-address</i> argument. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RIB quarantining detects mutually recursive routes and quarantines the last route that actually completes the mutual recursion. The quarantined route is periodically evaluated to see if the mutual recursion has gone away. If the recursion still exists, the route remains quarantined. If the recursion has gone away, the route is released from quarantine.

Use the **show route quarantined** command to display mutually recursive (looping) routes.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route quarantined** command:

```
RP/0/RSP0/CPU0:routerr# show route quarantined

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
       U - per-user static route, o - ODR, L - local

S    10.10.109.1/32 [1/0] via 10.10.34.1, 00:00:01 (quarantined)
      [1/0] via 10.10.37.1, 00:00:01 (quarantined)
      [1/0] via 10.10.60.1, 00:00:01 (quarantined)
      [1/0] via 10.10.68.1, 00:00:01 (quarantined)
      [1/0] via 10.10.91.1, 00:00:01 (quarantined)
      [1/0] via 10.10.93.1, 00:00:01 (quarantined)
      [1/0] via 10.10.97.1, 00:00:01 (quarantined)
S    10.0.0.0/8 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.0.0/16 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.10.0/24 [1/0] via 11.11.11.11, 00:01:29 (quarantined)
S    10.10.10.10/32 [1/0] via 11.11.11.11, 00:00:09 (quarantined)
```

This table describes the significant fields shown in the display.

Table 160: show route quarantined Field Descriptions

| Field | Description |
|----------------|--------------------------------------|
| 10.10.109.1/32 | IP address and length of the route. |
| [1/0] | Distance and metric for the route. |
| via 10.10.34.1 | IP address of next-hop on the route. |

| Field | Description |
|---------------|--|
| 00:00:01 | Time (in hh:mm:ss or <i>nd/h</i>) since the route was installed in the RIB. |
| (quarantined) | Shows that the route is quarantined. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current contents of the routing table. |

show route resolving-next-hop

To display the next-hop gateway or host to a destination address, use the **show route resolving-next-hop** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [ipv4| ipv6] [unicast| multicast| {topology topo-name}| safi-all]
resolving-next-hop ip-address [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| <i>ip-address</i> | IP address about which resolved next-hop information is to be displayed. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the next-hop gateway or host is displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route resolving-next-hop** command to perform a recursive route lookup on the supplied destination address and return information on the next immediate router (next hop) to the destination.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route resolving-next-hop** command:

```
RP/0/RSP0/CPU0:router# show route resolving-next-hop 10.1.1.1
NextHop matches 10.1.1.1/32
  Known via "local", distance 0, metric 0 (connected)
  Installed Aug 22 01:57:08.514
  Directly connected nexthops
    10.1.1.1 directly connected, via Loopback0
    Route metric is 0
```

This table describes the significant fields shown in the display.

Table 161: show route resolving-next-hop Field Descriptions

| Field | Description |
|-----------------|---|
| Known via | Name of the routing protocol that installed the matching route. |
| Route metric is | Metric of the route. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current contents of the routing table. |

show route static

To display the current static routes of the Routing Information Base (RIB), use the **show route static** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| {topology topo-name}| safi-all]
static [ standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the current static routes of the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route static** command to display information about static routes in the routing table.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route static** command:

```
RP/0/RSP0/CPU0:router# show route static
S    10.1.1.0/24 is directly connected, 00:54:05, GigabitEthernet3/0/0/1
S    192.168.99.99/32 [1/0] via 10.12.12.2, 00:54:04
```

This table describes the significant fields shown in the display.

Table 162: show route static Field Descriptions

| Field | Description |
|------------------------|--|
| S | Code to indicate the route is static. |
| 10.1.1.0/24 | IP address and distance for the route. |
| 00:54:05 | Time (in hh:mm:ss) since the route was installed in the RIB. |
| GigabitEthernet3/0/0/1 | Outbound interface for the route. |
| [1/0] | Distance and metric for the route. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current contents of the routing table. |

show route summary

To display the current contents of the Routing Information Base (RIB), use the **show route summary** command in EXEC mode.

```
show route [vrf {vrf-name| all}] [afi-all| ipv4| ipv6] [unicast| multicast| {topology topo-name}| safi-all]
summary [detail] [standby]
```

Syntax Description

| | |
|---|---|
| vrf { <i>vrf-name</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| afi-all | (Optional) Specifies all address families. |
| ipv4 | (Optional) Specifies IP Version 4 address prefixes. |
| ipv6 | (Optional) Specifies IP Version 6 address prefixes. |
| unicast | (Optional) Specifies unicast address prefixes. |
| multicast | (Optional) Specifies multicast address prefixes. |
| safi-all | (Optional) Specifies unicast and multicast address prefixes. |
| detail | (Optional) Displays a detailed summary of the contents of the RIB, including the number of paths and some protocol-specific route attributes. |
| standby | (Optional) Displays standby information. |

Command Default

If a **vrf** *vrf-name* is not specified, the contents of the RIB are displayed for the default IPv4 unicast VRF.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show route summary** command to display information about routes in the routing information base.

When a route summary is needed frequently—for instance, in a polling situation—use the **show route summary** command without the **detail** keyword. The **detail** keyword is used less frequently for verification purposes, because it is much more expensive (in bandwidth), requiring a scan of the entire routing database.

The **topology** keyword must be accompanied by the **ipv4 multicast** keywords, except when the **afi-all** keyword or the **safi-all** keyword is specified.

Task ID

| Task ID | Operations |
|---------|------------|
| rib | read |

Examples

The following is sample output from the **show route summary** command:

```
RP/0/RSP0/CPU0:router# show route summary

Route Source    Routes    Backup    Deleted    Memory (bytes)
static          1         0         0          136
connected       2         1         0          408
local           3         0         0          408
ospf            1673     2         0          272
isis            2         0         0          272
Total           10        1         0          1496
```

This table explains fields in the output of the **show route summary** command.

Table 163: show route summary Field Descriptions

| Field | Description |
|--------------|---|
| Route Source | Routing protocol name. |
| Routes | Number of selected routes that are present in the routing table for each route source. |
| Backup | Number of routes that are not selected (are backup to a selected route). |
| Deleted | Number of routes that have been marked for deletion in the RIB, but have not yet been purged. |
| Memory | Number of bytes allocated to maintain all routes for the particular route source. |

The following is sample output from the **show route summary** command with the **detail** keyword:

```
RP/0/RSP0/CPU0:router# show route summary detail

Route Source    Active Route    Active Path    Backup Route    Backup Path
static          1               1              0               0
connected       2               2              1               1
local           3               3              0               0
```

```

isis                1                1                1                1
Level 1:           0                0                1                1
Level 2:           1                1                0                0
ospf 1673          6                12               0                0
Intra-Area:        3                6                0                0
Inter-Area:        3                6                0                0
External-1:        0                0                0                0
External-2:        0                0                0                0
bgp 100            10               20               4                8
External:          5                10               4                8
Internal:          5                10               0                0
local:             0                0                0                0
Total              7                7                2                2

```

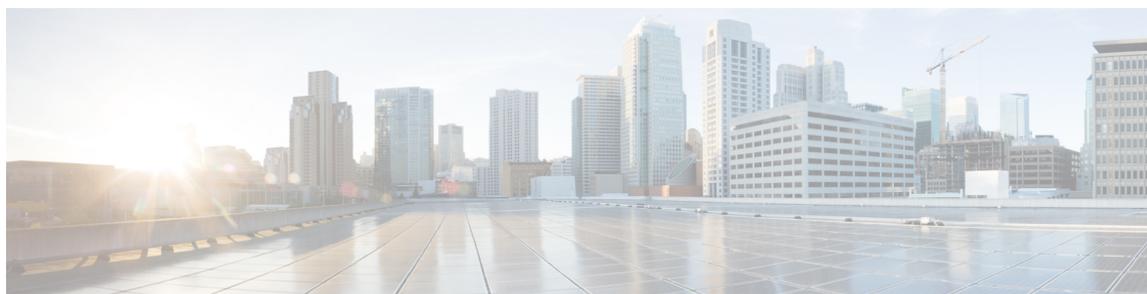
This table explains fields in the output of the **show route summary detail** command.

Table 164: show route summary detail Field Descriptions

| Field | Description |
|--------------|---|
| Route Source | Source of the route. Routing protocol name and type. |
| Active Route | Number of active routes present in the routing table for each route source. |
| Active Path | Number of active paths present in the routing table for each route source. |
| Backup Route | Number of routes that are backup to a selected route for each route source. |
| Backup Path | Number of paths that are backup to a selected path for each route source. |

Related Commands

| Command | Description |
|---|---|
| show route , on page 1412 | Displays the current contents of the routing table. |



RIP Commands

This module describes the commands used to configure and monitor the Routing Information Protocol (RIP).

For detailed information about RIP concepts, configuration tasks, and examples, see the *Implementing RIP on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [authentication keychain mode \(RIP\)](#), page 1445
- [auto-summary \(RIP\)](#), page 1447
- [broadcast-for-v2](#), page 1449
- [clear rip](#), page 1450
- [clear rip database](#), page 1452
- [clear rip interface](#), page 1454
- [clear rip out-of-memory](#), page 1456
- [clear rip statistics](#), page 1458
- [default-information originate \(RIP\)](#), page 1460
- [default-metric \(RIP\)](#), page 1462
- [distance \(RIP\)](#), page 1464
- [interface \(RIP\)](#), page 1467
- [maximum-paths \(RIP\)](#), page 1469
- [metric-zero-accept](#), page 1470
- [neighbor \(RIP\)](#), page 1471
- [nsf \(RIP\)](#), page 1473
- [output-delay](#), page 1474
- [passive-interface \(RIP\)](#), page 1475
- [poison-reverse](#), page 1477
- [receive version](#), page 1479
- [redistribute \(RIP\)](#), page 1481

- [router rip](#), page 1485
- [route-policy \(RIP\)](#), page 1487
- [send version](#), page 1489
- [show protocols \(RIP\)](#), page 1491
- [show rip](#), page 1493
- [show rip database](#), page 1495
- [show rip interface](#), page 1497
- [show rip statistics](#), page 1504
- [site-of-origin \(RIP\)](#), page 1506
- [split-horizon disable \(RIP\)](#), page 1508
- [timers basic](#), page 1510
- [validate-update-source disable](#), page 1512
- [vrf \(RIP\)](#), page 1513

authentication keychain mode (RIP)

To enable an authentication keychain mechanism on RIP interfaces, use the **authentication keychain mode** command in interface configuration mode or VRF-interface configuration mode. To disable authentication keychain configuration on RIP interfaces, use the **no** form of this command.

```
authentication keychain keychain_name mode {md5| text}
```

```
no authentication keychain keychain_name mode {md5| text}
```

Syntax Description

| | |
|----------------------|--|
| <i>keychain-name</i> | Specifies the name of the keychain configured using the keychain command. Note All keychains need to be configured in Cisco IOS XR keychain database using the keychain configuration commands described in <i>Implementing Keychain Management</i> module of <i>Cisco ASR 9000 Series Aggregation Services Router System Security Configuration Guide</i> |
| md5 | Specifies that the authentication keychain mode is keyed message digest (md5). |
| text | Specifies that the authentication keychain mode is clear text. |

Command Default

Keychain authentication is disabled.

Command Modes

Interface configuration
VRF-interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.0.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

All keychains need to be configured in Cisco IOS XR keychain database using the keychain configuration commands described in *Implementing Keychain Management* module of *Cisco ASR 9000 Series Aggregation Services Router System Security Configuration Guide*

Task ID

| Task ID | Operation |
|---------|-------------|
| rip | read, write |

Examples

This example shows how to configure an authentication keychain in md5 mode on a RIP VRF interface:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router rip
RP/0/RSP0/CPU0:router(config-rip)#vrf vrf_rip_auth
RP/0/RSP0/CPU0:router(config-rip-vrf)#interface POS 0/6/0/0
RP/0/RSP0/CPU0:router(config-rip-vrf-if)#authentication keychain key1 mode md5
```

This example shows how to configure an authentication keychain in clear text mode on a RIP interface:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router rip
RP/0/RSP0/CPU0:router(config-rip)#interface POS 0/6/0/1
RP/0/RSP0/CPU0:router(config-rip-if)#authentication keychain key2 mode text
```

Related Commands

| Command | Description |
|---|--|
| router rip , on page 1485 | Configures a routing process and enters router configuration mode for a Routing Information Protocol (RIP) process. |
| vrf (RIP) , on page 1513 | Defines a VPN routing and forwarding (VRF) instance and enters VRF configuration mode. Refer <i>Cisco ASR 9000 Series Aggregation Services Router System Security Command Reference</i> for complete command reference information. |
| key chain (key chain) | Creates or modifies a keychain. Refer <i>Cisco ASR 9000 Series Aggregation Services Router System Security Command Reference</i> for complete command reference information. |
| key (key chain) | Creates or modifies a keychain key. Refer <i>Cisco ASR 9000 Series Aggregation Services Router System Security Command Reference</i> for complete command reference information. |
| key-string (keychain) | Specifies text string for the key. Refer <i>Cisco ASR 9000 Series Aggregation Services Router System Security Command Reference</i> for complete command reference information. |

auto-summary (RIP)

To enable the automatic summarization of subnet routes into network-level routes, use the **auto-summary** command in the appropriate configuration mode. To disable this function and send subprefix routing information across classful network boundaries, use the **no** form of this command.

auto-summary

no auto-summary

Syntax Description This command has no arguments or keywords.

Command Default Disabled

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **auto-summary** command to turn on route summarization. Route summarization reduces the amount of routing information in the routing tables.

Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is off, subnets are advertised. Automatic summarization is disabled by default.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rip | read, write |

Examples The following example shows how to turn on RIP auto-summarization:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# auto-summary
```

Related Commands

| Command | Description |
|--|---|
| vrf (RIP) , on page 1513 | Defines a VRF instance and enters VRF configuration mode. |

broadcast-for-v2

To send Routing Information Protocol (RIP) Version 2 output packets to a broadcast address, use the **broadcast-for-v2** command in the appropriate configuration mode. To disable this feature, use the **no** form of this command.

broadcast-for-v2

no broadcast-for-v2

Syntax Description This command has no arguments or keywords.

Command Default RIPv2 output packets are not broadcasted.

Command Modes

- Router configuration
- VRF configuration
- Interface configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **broadcast-for-v2** command to broadcast RIP Version 2 broadcast updates to hosts that do not listen to multicasts. Version 2 updates (requests and responses) will be sent to the IP broadcast address 255.255.255.255 instead of the IP multicast address 244.0.0.9.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rip | read, write |

Examples The following example shows how to send RIP v2 output messages to a broadcast address for all RIP interfaces:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# broadcast-for-v2
```

clear rip

To clear VRF and interface-related information for Routing Information Protocol (RIP) such as database entries and statistics, use the **clear rip** command in EXEC configuration mode.

```
clear rip [vrf {vrf | all}]
```

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
|--|---|

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RIP VRFs and interfaces that are forcibly deactivated by the software because of a severe memory state are not activated again until the out-of-memory state is cleared by using the **clear rip**, **clear rip interface**, or **clear rip out-of-memory** command.

Task ID

| Task ID | Operations |
|---------|------------|
| rip | read |

Examples

The following example shows how to clear all database, interface, and VRF entries in RIP:

```
RP/0/RSP0/CPU0:router# clear rip vrf all
```

Related Commands

| Command | Description |
|---|---|
| clear rip interface, on page 1454 | Clears interface-related information for RIP such as database entries and statistics. |
| clear rip out-of-memory, on page 1456 | Clears the out-of-memory state for RIP. |

clear rip database

To clear only database entries from the Routing Information Protocol (RIP) topology table, use the **clear rip database** command in EXEC configuration mode.

clear rip [*vrf {vrf | all}*] **database** [*interface type interface-path-id*]

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| interface | (Optional) Specifies the interface to clear topology entries. |
| type | Interface type. For more information, use the question mark (?) online help function. |
| interface-path-id | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to clear only database entries from the topology table for the GigabitEthernet 0/1/0/0 interface:

```
RP/0/RSP0/CPU0:router# clear rip database interface GigabitEthernet 0/1/0/0
```

Related Commands

| Command | Description |
|---|--|
| show rip statistics, on page 1504 | Displays database and interface entry information for the RIP process. |

clear rip interface

To clear interface-related information for Routing Information Protocol (RIP) such as database entries and statistics, use the **clear rip interface** command in EXEC configuration mode.

clear rip [*vrf* {*vrf* | **all**}] **interface** *type interface-path-id*

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| interface | Specifies the interface to clear topology entries. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RIP VRFs and interfaces that are forcibly deactivated by the software because of a severe memory state are not activated again until after out-of-memory state is cleared by using the **clear rip** , **clear rip interface** or **clear rip out-of-memory** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to clear all interface-related data such as routes and statistics from the GigabitEthernet 0/1/0/0 interface:

```
RP/0/RSP0/CPU0:router# clear rip vrf vpn-1 interface GigabitEthernet 0/1/0/0
```

Related Commands

| Command | Description |
|---|---|
| clear rip, on page 1450 | Clears VRF and interface-related information for RIP such as database entries and statistics. |
| clear rip out-of-memory, on page 1456 | Clears the out-of-memory state for RIP. |

clear rip out-of-memory

To clear the out-of-memory state for Routing Information Protocol (RIP), use the **clear rip out-of-memory** command in EXEC configuration mode.

clear rip [*vrf* {*vrf* | **all**}] **out-of-memory** [**interface** *type interface-path-id*]

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| interface | (Optional) Specifies the interface to clear topology entries. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **clear rip out-of-memory** command, to clear the out-of-memory state completely and allow the RIP process to force the VRF or interface to shut down.

When the router begins to run out of memory, the RIP process can transition through different memory states defined as Normal, Minor, Severe, and Critical:

- In Normal state, RIP VRFs and interfaces function normally.

- In Minor state, RIP VRFs and interfaces that are currently active are allowed to remain active. VRFs and interfaces that are not currently active are not allowed to become active until the RIP process transitions to Normal state.
- In Severe state, a few VRFs and interfaces are forcibly brought down periodically until the RIP process transitions to another state.
- In Critical state, the RIP process is forcibly shut down.

VRFs and interfaces that are forcibly shut down in Severe state are not automatically activated when the RIP process transitions to Minor or Normal state. When a VRF or interface is forcibly brought down while in Severe state, the **clear rip**, **clear rip interface** or **clear rip out-of-memory** command clears the Forced Down state and reactivates the VRF or interface.

The **show rip** and **show rip interface** commands allow you to view the current out-of-memory state.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to clear the out-of-memory state for a RIP process:

```
RP/0/RSP0/CPU0:router# clear rip out-of-memory
```

Related Commands

| | |
|---|---|
| clear rip, on page 1450 | Clears VRF and interface-related information for RIP such as database entries and statistics. |
| clear rip interface, on page 1454 | Clears interface-related information for RIP such as database entries and statistics. |
| show rip, on page 1493 | Displays configuration and status of RIP. |
| show rip interface, on page 1497 | Displays interface entry information from the RIP topology table. |

clear rip statistics

To clear the Routing Information Protocol (RIP) statistics, use the **clear rip statistics** command in EXEC configuration mode.

clear rip [**vrf** {*vrf* | **all**}] **statistics** [**interface** *type interface-path-id*]

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| interface | (Optional) Specifies the interface from which to clear topology entries. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

EXEC configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to clear all RIP statistics:

```
RP/0/RSP0/CPU0:router# clear rip statistics
```

Related Commands

| Command | Description |
|---|--|
| show rip statistics, on page 1504 | Displays database and interface entry information for the RIP process. |

default-information originate (RIP)

To generate a default route into Routing Information Protocol (RIP), use the **default-information originate** command in the appropriate configuration mode. To disable a default route into RIP, use the **no** form of this command.

default-information originate [*route-policy name*]

no default-information originate

Syntax Description

| | |
|---------------------------------|--|
| route-policy <i>name</i> | Route policy name that indicates criteria for the default route. |
|---------------------------------|--|

Command Default

This command is disabled by default.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to originate a default route in RIP updates based on the result of running the route policy on the routing table:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# default-information originate route-policy policy1
```

Related Commands

| Command | Description |
|---|--|
| route-policy (RIP) , on page 1487 | Applies a routing policy to updates advertised to or received from a RIP neighbor. |

default-metric (RIP)

To set default metric values for routes redistributed from other protocols into Routing Information Protocol (RIP), use the **default-metric** command in the appropriate configuration mode. To return to the default state, use the **no** form of this command.

default-metric *number-value*

no default-metric

Syntax Description

| | |
|---------------------|---|
| <i>number-value</i> | Default metric value. Range is 1 to 15. |
|---------------------|---|

Command Default

Default metrics are not set.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **default-metric** command with the **redistribute** command to cause RIP to use the same metric value for all redistributed routes. A default metric helps solve the problem of redistributing routes with incompatible metrics by providing a reasonable substitute and enables redistribution to proceed. If you want to set different metrics for other redistributed protocols, use the **route-policy** option in the **redistribute** command.

The RIP metric used for redistributed routes is determined by the route policy. If a route policy is not configured or the route policy does not set the RIP metric, the metric is determined based on the redistributed protocol. For VPNv4 routes redistributed by BGP, the RIP metric set at the remote PE router is used, if valid.

In all other cases (BGP, IS-IS, OSPF, EIGRP, connected, static), the metric set by the **default-metric** command is used. If a valid metric cannot be determined, then redistribution does not happen.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how a router in autonomous system 109 uses both the RIP and the Open Shortest Path First (OSPF) routing protocols. The example advertises OSPF-derived routes using RIP and assigns the OSPF-derived routes a RIP metric of 10:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# default-metric 10
RP/0/RSP0/CPU0:router(config-rip-vrf)# redistribute ospf 109
```

Related Commands

| Command | Description |
|--|--|
| redistribute (RIP), on page 1481 | Redistributes routes from one routing domain into RIP. |

distance (RIP)

To define the administrative distance assigned to routes discovered by the Routing Information Protocol (RIP), use the **distance admin-distance** command in the appropriate configuration mode. To remove the distance definition from the configuration file and restore the system to its default condition, use the **no** form of this command.

distance *admin-distance* [*prefix prefix-length*] *prefix mask*

no distance *admin-distance*

Syntax Description

| | |
|-----------------------|--|
| <i>admin-distance</i> | Administrative distance to be assigned to RIP routes. Range is 0 to 255. |
| <i>prefix</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>prefix-length</i> | (Optional) The <i>prefix-length</i> argument specifies the length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash must precede the decimal value. Range is 0 to 32 for IPv4 addresses. |
| <i>mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 shows that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address. |

Command Default

admin-distance : 120

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **distance** command to change the preference of RIP routes over other protocol routes. When administrative distance and redistribution features are used together, routing behavior may be affected for routes accepted from and advertised to RIP neighbors.

Numerically, an administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored.

The order in which you enter distance commands can affect the assigned administrative distances in unexpected ways.

This table lists default administrative distances.

Table 165: Default Administrative Distances of Routing Protocols

| Routing Protocol | Administrative Distance Value |
|-------------------------------|-------------------------------|
| Connected interface | 0 |
| Static route out an interface | 0 |
| Static route to next-hop | 1 |
| EIGRP Summary Route | 5 |
| External BGP | 20 |
| Internal EIGRP | 90 |
| OSPF | 110 |
| IS-IS | 115 |
| RIP Versions 1 and 2 | 120 |
| External EIGRP | 170 |
| Internal BGP | 200 |
| Unknown | 255 |

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to set the administrative distance for a particular prefix:

```
RP/0/RSP0/CPU0:router(config)# router rip
```

distance (RIP)

```
RP/0/RSP0/CPU0:router(config-rip)# distance 85 192.168.10.0/24
```

Related Commands

| Command | Description |
|--|--|
| redistribute (RIP), on page 1481 | Redistributes routes from one routing domain into RIP. |

interface (RIP)

To define the interfaces on which the Routing Information Protocol (RIP) runs and enter interface configuration mode, use the **interface** command in router configuration mode. To disable RIP routing for interfaces, use the **no interface** form of this command.

interface *type interface-path-id*

no interface *type interface-path-id*

Syntax Description

| | |
|--------------------------|--|
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or a virtual :interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

When you do not specify this command in configuration mode, RIP routing for interfaces is not enabled.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **interface** command to associate a specific interface with a RIP process. The interface remains associated with the process even when the IPv4 address of the interface changes.

This command places the router in interface configuration mode, from which you can configure interface-specific settings. Commands configured under this mode (such as the [broadcast-for-v2](#), on page 1449 command) are automatically bound to that interface.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to enter interface configuration mode for a RIP process, and send RIP Version 2 messages to the broadcast address on the GigabitEthernet interface 0/1/0/0:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# ?

broadcast-for-v2    Specify broadcast address for RIP v2 output packet
commit             Commit the configuration changes to running
describe           Describe a command without taking real actions
do                Run an exec command
exit              Exit from this submode
metric-zero-accept Accept rip update with metric 0 to compensate a common bug
no                Negate a command or set its defaults
passive-interface  Suppress routing updates on this interface
poison-reverse     Enable poison reverse
receive           Advertisement reception
route-policy       Apply route policy to routing updates
send              Advertisement transmission
show              Show contents of configuration
site-of-origin     SOO community for prefixes learned over this interface
split-horizon     Disable split horizon
RP/0/RSP0/CPU0:router(config-rip-if)# broadcast-for-v2
```

Related Commands

| Command | Description |
|--|--|
| broadcast-for-v2, on page 1449 | Sends RIP Version 2 output packets to a broadcast address. |

maximum-paths (RIP)

To configure the maximum number of equal cost parallel routes that the Routing Information Protocol (RIP) will install into the routing table, use the **maximum-paths** command in the appropriate configuration mode. To remove the **maximum-paths** command from the configuration file and restore the system to its default condition with respect to RIP, use the **no** form of this command.

maximum-paths *maximum*

no maximum-paths

| Syntax Description | maximum | Maximum number of parallel routes that RIP can install in a routing table. Range is 1 to 32.. |
|--------------------|---------|---|
|--------------------|---------|---|

Command Default 4 paths

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to allow a maximum of 16 equal cost paths to a destination:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# maximum-paths 16
```

metric-zero-accept

To allow RIP to accept routing entries from RIP updates with a metric set to zero (0), use the **metric-zero-accept** command in interface configuration mode. To remove the **metric-zero-accept** command from the configuration file and restore the system to its default condition with respect to RIP, use the **no** form of this command.

metric-zero-accept

no metric-zero-accept

Syntax Description This command has no arguments or keywords.

Command Default RIP routes received with a metric of zero (0) are ignored.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

After the **metric-zero-accept** command is configured on routing entries from RIP updates, RIP accepts these routes and then sets the metric to one (1).

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to set the RIP interface to accept metric zero on routing entries:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# metric-zero-accept
```

neighbor (RIP)

To define a neighboring router with which to exchange Routing Information Protocol (RIP) information, use the **neighbor** command in the appropriate configuration mode. To remove an entry, use the **no** form of this command.

neighbor *ip-address*

no neighbor *ip-address*

| Syntax Description | |
|--------------------|--|
| ip-address | IP address of a peer router with which routing information is exchanged. |

| Command Default | No neighboring routers are defined. |
|-----------------|-------------------------------------|
|-----------------|-------------------------------------|

| Command Modes | Router configuration VRF configuration |
|---------------|---|
|---------------|---|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

| Usage Guidelines | To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance. |
|------------------|---|
|------------------|---|

Use the **neighbor** command to permit the point-to-point (nonbroadcast) exchange of routing information. When the **neighbor** command is used in combination with the **passive-interface** command in router configuration mode, routing information can be exchanged between a subset of routers and access servers on a LAN.

Multiple **neighbor** commands can be used to specify additional neighbors or peers.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples

The following example shows how to permit the sending of RIP updates to specific neighbors. One copy of the update is generated per neighbor:

```
RP/0/RSP0/CPU0:router(config)# router rip  
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1  
RP/0/RSP0/CPU0:router(config-rip-vrf)# neighbor 172.16.1.2
```

Related Commands

| Command | Description |
|--|--|
| passive-interface (RIP) , on page 1475 | Suppresses the sending of RIP updates on an interface. |

nsf (RIP)

To configure nonstop forwarding (NSF) on Routing Information Protocol (RIP) routes after a RIP process shutdown or restart, use the **nsf** command in the appropriate configuration mode. To remove this command from the configuration file and restore the system to its default condition, use the **no** form of this command.

nsf

no nsf

Syntax Description This command has no arguments or keywords.

Command Default NSF is disabled.

Command Modes Router configuration
VRF configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When you use the **nsf** command, NSF lifetime is automatically set to two times the update time (with a minimum value of 60 seconds). The RIP process must reconverge within this time. If the convergence exceeds the NSF lifetime, routes are purged from the Routing Information Base (RIB) and NSF may fail.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to configure RIP NSF:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# nsf
```

output-delay

To change the interpacket delay for Routing Information Protocol (RIP) updates sent, use the **output-delay** command in the appropriate configuration mode. To remove the delay, use the **no** form of this command.

output-delay *delay*

no output-delay *delay*

Syntax Description

| | |
|-------|---|
| delay | Delay (in milliseconds) between consecutive packets in a multiple-packet RIP update. The range is from 8 to 50. |
|-------|---|

Command Default

The default is no delay.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **output-delay** command if you are sending at high speed to a low-speed router that might not be able to receive at the high speed. Configuring this command helps prevent the routing table from losing information.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to set the interpacket delay to 10 milliseconds:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# output-delay 10
```

passive-interface (RIP)

To suppress the sending of Routing Information Protocol (RIP) updates on an interface, use the **passive-interface** command in interface configuration mode. To unsuppress updates, use the **no** form of this command.

passive-interface

no passive-interface

Syntax Description This command has no arguments or keywords.

Command Default RIP updates are sent on the interface.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

While RIP stops sending routing updates to the multicast (or broadcast) address on a passive interface, RIP continues to receive and process routing updates from its neighbors on that interface.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows that GigabitEthernet interface 0/1/0/0 stops multicasting (or broadcasting) RIP updates while continuing to receive RIP updates normally. GigabitEthernet interface 0/1/0/3 sends and receives updates normally. Also RIP updates are unicast to neighbor 172.168.1.2 over the appropriate interface:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# neighbor 172.16.1.2
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# passive-interface
RP/0/RSP0/CPU0:router(config-rip-if)# exit
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/3
RP/0/RSP0/CPU0:router(config-rip-if)# exit
```

Related Commands

| Command | Description |
|--|---|
| neighbor (RIP), on page 1471 | Defines a neighboring router with which to exchange RIP protocol information. |

poison-reverse

To enable poison reverse processing of Routing Information Protocol (RIP) router updates, use the **poison-reverse** command in interface configuration mode. To disable poison reverse processing of RIP updates, use the **no** form of this command.

poison-reverse

no poison-reverse

Syntax Description This command has no arguments or keywords.

Command Default Poison reverse processing is disabled.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Route poisoning prevents routing loops by communicating to other routers that a route is no longer reachable, effectively removing these routes from other router's routing tables. The system default, **split horizon**, provides that routes learned through RIP are not advertised from the interface over which they were learned.

The **poison-reverse** command enables poison reverse processing of RIP router updates. A router that receives route poisoning information sends the poisoning information back to the sending router, a process called poison reverse. This process ensures that all routers on the same interface have received the poisoned route information.

If both **poison-** reverse and split horizon are configured, then simple split horizon behavior (suppression of routes from the interface over which they were learned) is replaced by poison reverse behavior. If split horizon is disabled, the poison reverse configuration is ignored.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples

The following example shows how to enable poison reverse processing for an interface running RIP:

```
RP/0/RSP0/CPU0:router(config)# router rip  
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-rip-if)# poison-reverse
```

Related Commands

| Command | Description |
|---|---------------------------------------|
| split-horizon disable (RIP), on page 1508 | Disables the split horizon mechanism. |

receive version

To configure the Routing Information Protocol (RIP) interface to accept version-specific packets, use the **receive version** command in interface configuration mode. To revert to the default setting, use the **no** form of this command.

receive version {1|2|1 2}

no receive version {1|2|1 2}

Syntax Description

| | |
|-----|--------------------------------|
| 1 | Version 1 packets. |
| 2 | Version 2 packets. |
| 1 2 | Both versions 1 and 2 packets. |

Command Default

Version 2

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **receive version** command to override the default behavior of RIP. This command applies only to the interface being configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to configure an interface to accept both RIP Version 1 and 2 packets:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# receive version 1 2
```

Related Commands

| Command | Description |
|--|--|
| send version, on page 1489 | Configures the RIP interface to send version specific packets. |

redistribute (RIP)

To redistribute routes from another routing domain into Routing Information Protocol (RIP), use the **redistribute** command in the appropriate configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

Border Gateway Protocol (BGP)

```
redistribute bgp process-id [route-policy name] [external| internal| local]
```

```
no redistribute bgp process-id
```

Connected Interface Routes

```
redistribute connected [route-policy name]
```

```
no redistribute connected
```

Enhanced Interior Gateway Routing Protocol (EIGRP)

```
redistribute eigrp process-id [route-policy name]
```

```
no redistribute eigrp process-id
```

Intermediate System-to-Intermediate System (ISIS)

```
redistribute isis process-id [route-policy name] [level-1| level-1-2| level-2]
```

```
no redistribute isis process-id
```

Open Shortest Path First (OSPF)

```
redistribute ospf process-id [route-policy name] [match {external [1| 2]| internal| nssa-external [1| 2]}]
```

```
no redistribute ospf process-id
```

IP Static Routes

```
redistribute static [route-policy name]
```

```
no redistribute static
```

Syntax Description

| | |
|-----|---|
| bgp | Distributes routes from the BGP protocol. |
|-----|---|

| | |
|---------------------------------|--|
| process-id | <ul style="list-style-type: none"> • For the bgp keyword: <ul style="list-style-type: none"> ◦ Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. ◦ Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. • For the eigrp keyword, an EIGRP instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. • For the isis keyword, an IS-IS instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. • For the ospf keyword, an OSPF instance name from which routes are to be redistributed. The value takes the form of a string. A decimal number can be entered, but it is stored internally as a string. |
| external | (Optional) Specifies BGP external routes only. |
| internal | (Optional) Specifies BGP internal routes only. |
| local | (Optional) Specifies BGP local routes only. |
| route-policy <i>name</i> | (Optional) Specifies the identifier of a configured policy. A policy is used to filter the importation of routes from this source routing protocol to RIP. |
| level-1 | (Optional) Redistributes Level 1 IS-IS routes into other routing protocols independently. |
| level-1-2 | (Optional) Distributes both Level 1 and Level 2 IS-IS routes into other routing protocols. |
| level-2 | (Optional) Distributes Level 2 IS-IS routes into other routing protocols independently. |

[**match** { **external** [**1** | **2**] | **internal** | **nssa-external** [**1** | **2**] } [**route-policy** *name*]

(Optional) Specifies the criteria by which OSPF routes are redistributed into other routing domains. It can be one or more of the following:

- **internal** —Routes that are internal to a specific autonomous system (intra- and inter-area OSPF routes).
- **external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 external routes.
- **nssa-external** [**1** | **2**]—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 or Type 2 not-so-stubby area (NSSA) external routes.

For the external and nssa-external options, if a type is not specified, then both Type 1 and Type 2 are assumed.

If no match is specified, the default is no filtering.

static Redistributes IP static routes.

Command Default

Route redistribution is disabled.

Command Modes

Router configuration
VRF configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

When you are redistributing routes (into RIP) using both command keywords for setting or matching of attributes and a route policy, the routes are run through the route policy first, followed by the keyword matching and setting.

Redistributed routing information may be filtered by the **route-policy** *name* keyword and argument. This filtering ensures that only those routes intended by the administrator are redistributed by RIP.

The RIP metric used for redistributed routes is determined by the route policy. If a route policy is not configured or the route policy does not set the RIP metric, the metric is determined based on the redistributed protocol. For VPNv4 routes redistributed by BGP, the RIP metric set at the remote PE router is used, if valid.

In all other cases (BGP, IS-IS, OSPF, EIGRP, connected, static), the metric set by the **default-metric** command is used. If a valid metric cannot be determined, then redistribution does not happen.

For information about routing policies, see the *Routing Policy Commands on Cisco ASR 9000 Series Router* module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Command Reference*.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to cause BGP routes to be redistributed into a RIP process:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# redistribute bgp 100
```

Related Commands

| Command | Description |
|--|--|
| default-metric (RIP), on page 1462 | Sets default metric values for routes redistributed from other protocols into RIP. |

router rip

To configure a routing process and enter router configuration mode for a Routing Information Protocol (RIP) process, use the **router rip** command in global configuration mode. To turn off the RIP routing process, use the **no** form of this command.

router rip

no router rip

Syntax Description This command has no arguments or keywords.

Command Default No router process is defined.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to configure a router process for RIP:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)#
```

The following example shows how to enter router configuration mode for RIP and identify commands that can be issued from that mode.

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# ?

auto-summary          Enable automatic network number summarization
broadcast-for-v2      Send RIP v2 output packets to broadcast address
commit                Commit the configuration changes to running
```

```
default-information      Control distribution of default information
default-metric          Set metric of redistributed routes
describe                Describe a command without taking real actions
distance                Define an administrative distance
do                      Run an exec command
exit                    Exit from this submode
interface                Enter the RIP interface configuration submode
maximum-paths           Maximum number of paths allowed per route
neighbor                Specify a neighbor router
no                      Negate a command or set its defaults
nsf                     Enable Cisco Non Stop Forwarding
output-delay            Interpacket delay for RIP updates
redistribute             Redistribute information from another routing protocol
route-policy            Apply route policy to routing updates
show                    Show contents of configuration
timers                  Adjust routing timers
validate-update-source  Validate source address of routing updates
vrf                     Enter the RIP vrf configuration submode
RP/0/RSP0/CPU0:router(config-rip)#
```

route-policy (RIP)

To apply a routing policy to updates advertised to or received from a Routing Information Protocol (RIP) neighbor, use the **route-policy** command in the appropriate configuration mode. To disable applying routing policy to updates, use the **no** form of this command.

route-policy *name* {**in**|**out**}

no route-policy *name* {**in**|**out**}

Syntax Description

| | |
|------|------------------------------------|
| name | Name of route policy. |
| in | Applies policy to inbound routes. |
| out | Applies policy to outbound routes. |

Command Default

No policy is applied.

Command Modes

Router configuration
VRF configuration
Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to specify a routing policy for an inbound or outbound route. The policy can be used to filter routes or modify route attributes.



Note

If a route policy is configured both on the interface and on the VRF, the interface route policy is applied.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to filter routing updates received on an interface:

```
RP/0/RSP0/CPU0:router(config)# router rip  
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-rip-if)# route-policy updpol-1 in
```

send version

To configure the Routing Information Protocol (RIP) interface to send version specific packets, use the **send version** command in interface configuration mode. To revert to the default setting, use the **no** form of this command.

send version {1| 2| 1 2}

no send version {1| 2| 1 2}

Syntax Description

| | |
|-----|---------------------------------------|
| 1 | Version 1 packets. |
| 2 | Version 2 packets. |
| 1 2 | Both Version 1 and Version 2 packets. |

Command Default

Version 2

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **send version** command to override the default behavior of RIP. This command applies only to the interface being configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to configure an interface to send only RIP Version 2 packets:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# send version 2
```

send version

Related Commands

| Command | Description |
|---|--|
| receive version, on page 1479 | Configures the RIP interface to accept version-specific packets. |

show protocols (RIP)

To display information about the Routing Information Protocol (RIP) process configuration, use the **show protocols** command in EXEC mode.

```
show protocols [ipv4|afi-all] [all|protocol] [default-context|vrf vrf-name] [private]
```

Syntax Description

| | |
|----------------------------|---|
| ipv4 | (Optional) Specifies an IPv4 address family. |
| afi-all | (Optional) Specifies all address families. |
| all | (Optional) Specifies all protocols for a given address family. |
| <i>protocol</i> | (Optional) Specifies a routing protocol. <ul style="list-style-type: none"> For the IPv4 address family, the options are eigrp, bgp, isis, ospf, and rip. |
| default-context | (Optional) Displays default context information. This keyword is available when the eigrp or rip protocol is specified. |
| vrf <i>vrf-name</i> | (Optional) Displays VPN routing and forwarding (VRF) information for the specified process. This keyword is available when the eigrp or rip protocol is specified. |
| private | (Optional) Displays private EIGRP data. This keyword is available when the eigrp protocol is specified. |

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show protocols** command to get information about the protocols running on the router and to quickly determine which protocols are active. The command summarizes the important characteristics of the running protocol, and command output varies depending on the specific protocol selected.

For RIP, the command output lists the instance number, default AS context, router ID, default networks, distance, maximum paths, and so on.

Task ID

| Task ID | Operations |
|---------|------------|
| RIP | read |

Examples

The following is sample output from the **show protocols rip** command:

```
RP/0/RSP0/CPU0:router# show protocols rip
Routing Protocol RIP
 2 VRFs (including default) configured, 2 active
 25 routes, 16 paths have been allocated
Current OOM state is "Normal"
UDP socket descriptor is 37
  VRF      Active  If-config  If-active  Routes    Paths    Updates
  default  Active   3          3          11        7        30s
```

This table describes the significant fields shown in the display.

Table 166: show protocols Field Descriptions

| Field | Description |
|-----------------|---|
| VRFs configured | Number of VRFs configured. |
| VRFs active | Number of active VRFs. |
| Routes | Number of allocated routes. |
| Paths | Number of allocated paths. |
| OOM state | Current out-of-memory state of RIP process. |
| UDP socket | Current UDP socket descriptor value. |

show rip

To display configuration and status of Routing Information Protocol (RIP), use the **show rip** command in EXEC mode.

```
show rip [vrf {vrf | all}]
```

| | | |
|---------------------------|--|---|
| Syntax Description | vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
|---------------------------|--|---|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|------|
| Command Modes | EXEC |
|----------------------|------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rip | read |

Examples

The following example shows sample output from the **show rip** command:

```
RP/0/RSP0/CPU0:router# show rip

RIP config:
Active?:                               Yes
Added to socket?:                       Yes
Out-of-memory state:                    Normal
Version:                                 2
Default metric:                          Not set
Maximum paths:                           4
Auto summarize?:                         No
Broadcast for V2?:                       No
Packet source validation?:              Yes
NSF:                                     Disabled
Timers: Update:                          30 seconds (25 seconds until next update)
```

```

Invalid:          180 seconds
Holddown:        180 seconds
Flush:           240 seconds

```

This table describes the significant fields shown in the display.

Table 167: show rip Field Descriptions

| Field | Description |
|---------------------------|---|
| Active? | Active state setting. |
| Added to socket? | Multicast group setting on RIP configured interfaces. If yes, updates are received on these interfaces. |
| Out-of-memory state | Out-of-memory state for RIP can be one of the following: Normal, Minor, Severe, or Critical. |
| Version | Version number is 2. |
| Default metric | Default metric value, if configured. Otherwise Not set. |
| Maximum paths | Number of maximum paths allowed per RIP route. |
| Auto summarize? | Auto-summarize state setting. |
| Broadcast for V2? | RIP Version 2 broadcast setting. |
| Packet source validation? | Validation setting for the source IP address of incoming routing updates to RIP. |
| Timers | RIP network timer settings. |

show rip database

To display database entry information from the Routing Information Protocol (RIP) topology table, use the **show rip database** command in EXEC mode.

```
show rip [vrf {vrf | all}] database [prefix prefix-length| prefix mask]
```

Syntax Description

| | |
|--|--|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| <i>prefix</i> | (Optional) Network IP address about which routing information should be displayed. |
| <i>prefix-length</i> | (Optional) The <i>prefix-length</i> argument specifies the length of the IP prefix. A decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix (the network portion of the address). A slash between must precede the decimal value. Range is 0 to 32 for IPv4 addresses. |
| <i>prefix-mask</i> | (Optional) Network mask specified in either of two ways: <ul style="list-style-type: none"> • Network mask can be a four-part, dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • Network mask can be indicated as a slash (/) and number. For example, /8 shows that the first 8 bits of the mask are ones, and the corresponding bits of the address are the network address. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Summary address entries appear in the database only if relevant child routes are summarized. When the last child route for a summary address becomes invalid, the summary address is also removed from the routing table.

Task ID

| Task ID | Operations |
|---------|------------|
| rip | read |

Examples

The following is sample output from the **show rip database** command:

```
RP/0/RSP0/CPU0:router# show rip database

Routes held in RIP's topology database:
10.0.0.0/24
  [0] directly connected, GigabitEthernet0/6/0/0
10.0.0.0/8 auto-summary
12.0.0.0/24
  [5] distance: 20 redistributed
12.0.0.0/8 auto-summary
50.50.0.0/24
  [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0
50.50.1.0/24 (inactive)
  [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0
50.0.0.0/8 auto-summary
90.90.0.0/24
  [5] distance: 20 redistributed
90.90.1.0/24
  [5] distance: 20 redistributed
```

This table describes the significant fields shown in the display.

Table 168: show rip database Field Descriptions

| Field | Description |
|--|--|
| 10.0.0.0/24 [0] directly connected, GigabitEthernet0/6/0/0 | Prefix and prefix length for a RIP connected route. 10.0.0.0/24 is directly connected to GigabitEthernet 0/6/0/0. The [0] represents the metric. |
| 10.0.0.0/8 auto-summary | 10.0.0.0/8 is a summary route entry. |
| 12.0.0.0/24 [5] distance: 20 redistributed | 12.0.0.0/24 is a redistributed route. The metric is 5, and the distance is 20. |
| 50.50.0.0/24 [1] via 10.0.0.20, next-hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0 | The destination route 50.50.0.0/24 is learned through RIP, and the source 10.0.0.20 advertised it from GigabitEthernet 0/6/0/0. The route was last updated one second ago. |
| 50.50.1.0/24 (inactive) [1] via 10.0.0.20, next hop 10.0.0.20, Uptime: 1s, GigabitEthernet0/6/0/0 | The destination route 50.50.1.0/24 is not active in the routing table. |

show rip interface

To display interface entry information from the Routing Information Protocol (RIP) topology table, use the **show rip interface** command in EXEC mode.

```
show rip [vrf {vrf | all}] interface [type interface-path-id]
```

Syntax Description

| | |
|--|--|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
| interface | (Optional) Specifies the interface from which to clear topology entries. |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|---|
| Release 4.0.0 | The command output was modified to include authentication keychain configuration information. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| rip | read |

Examples

This example is sample output from the **show rip interface** command:

```
RP/0/RSP0/CPU0:router# show rip interface

GigabitEthernet0_6_0_0
Rip enabled?:                Yes
Out-of-memory state:        Normal
Broadcast for V2:           No
Accept Metric 0?:          No
Send versions:              2
Receive versions:          2
Interface state:            Up
IP address:                 10.0.0.12/24
Metric Cost:                0
Split horizon:              Enabled
Poison Reverse:             Disabled
Joined multicast group?:    Yes

GigabitEthernet0_6_0_2
Out-of-memory state:        Normal
Rip enabled?:                Yes
Broadcast for V2:           No
Accept Metric 0?:          No
Send versions:              2
Receive versions:          2
Interface state:            Up
IP address:                 12.0.0.12/24
Metric Cost:                0
Split horizon:              Enabled
Poison Reverse:             Disabled
Joined multicast group?:    Yes

RIP peers attached to this interface:
 12.0.0.13
  uptime: 3    version: 2
  packets discarded: 0    routes discarded: 402
```

This table describes the significant fields shown in the display.

Table 169: show rip interface Field Descriptions

| Field | Description |
|---------------------|--|
| Rip enabled? | Specifies whether the RIP routing protocol is enabled on the interface. |
| Out-of-memory state | Specifies the current out-of-memory state on the interface. |
| Broadcast for V2 | Specifies whether RIP Version 2 output packets are sent to a broadcast address on the interface. |
| Accept Metric 0? | Specifies whether this interface accepts routing entries from RIP updates with a metric set to zero (0). |
| Send versions: | Specifies which version RIP uses to send out packets on this interface. |

| Field | Description |
|---|---|
| Receive versions: | Specifies which version packets RIP accepts on this interface. |
| Interface state: | Specifies whether the interface is in an up or a down state. |
| IP address: | IP address of the interface. |
| Metric Cost: | Specifies metric cost value. |
| Split horizon: | Specifies whether split horizon is enabled on this interface. |
| Poison Reverse: | Specifies whether poison reverse is enabled on this interface. |
| Joined multicast group?: | Specifies whether the interface has joined the RIP multicast group 224.0.0.9. |
| RIP peers attached to this interface 12.0.0.13 | List of RIP neighbors on this interface. |
| uptime: 3 | Specifies how long this neighbor is up. |
| version: 2 | Specifies which version packets are received from this neighbor. |
| packets discarded: 0 | Specifies the number of packets discarded from this neighbor. |
| routes discarded: 402 | Specifies the number of routes discarded from this neighbor. |

Authentication Keychain Configuration for RIP Interface on Default VRF

These examples are output of the **show rip interface *interface-path-id*** command to display authentication keychain configuration for RIP interface on default VRF.

When an existing keychain with MD5 cryptographic algorithm was configured on the RIP interface:

```
GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?:                No
Out-of-memory state:        Normal
Broadcast for V2:           No
Accept Metric 0?:          No
Send versions:              2
Receive versions:           2
Interface state:            Unknown State
IP address:                 0.0.0.0/0
Metric Cost:                0
Split horizon:              Enabled
```

```

Poison Reverse:           Disabled
Socket set options:
Joined multicast group?:   No
LPTS filter set?:         No
Authentication mode: MD5  Key chain: <key-chain-name>
Current active send key id: <send key id>
Current active receive key id: <recv key id>
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>

```

When the keychain configured on the RIP interface does not exist or does not have any active keys:

```

GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?:             No
Out-of-memory state:     Normal
Broadcast for V2:        No
Accept Metric 0?:        No
Send versions:           2
Receive versions:        2
Interface state:         Unknown State
IP address:               0.0.0.0/0
Metric Cost:              0
Split horizon:           Enabled
Poison Reverse:          Disabled
Socket set options:
Joined multicast group?:   No
LPTS filter set?:         No

Authentication mode: MD5  Key chain: <key-chain-name>
No active key found in keychain database.
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>

```

When an active key exists in the keychain configured on the RIP interface, but not configured with MD5 cryptographic algorithm:

```

GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?:             No
Out-of-memory state:     Normal
Broadcast for V2:        No
Accept Metric 0?:        No
Send versions:           2
Receive versions:        2
Interface state:         Unknown State
IP address:               0.0.0.0/0
Metric Cost:              0
Split horizon:           Enabled
Poison Reverse:          Disabled
Socket set options:
Joined multicast group?:   No
LPTS filter set?:         No

Authentication mode: MD5  Key chain: <key-chain-name>
Key(s) not configured with MD5 cryptographic algorithm.
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>

```

When no authentication keychain was configured on the RIP interface:

```
GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?: No
Out-of-memory state: Normal
Broadcast for V2: No
Accept Metric 0?: No
Send versions: 2
Receive versions: 2
Interface state: Unknown State
IP address: 0.0.0.0/0
Metric Cost: 0
Split horizon: Enabled
Poison Reverse: Disabled
Socket set options:
Joined multicast group?: No
LPTS filter set?: No
```

```
Authentication mode is not set.
Packets received: <num-rx-packets>
```

Authentication keychain Configuration for RIP Interface on Non-default VRF

These examples are output of the **show rip vrf vrf-name interface interface-path-id** command to display authentication keychain configuration for RIP interface on a non- default VRF.

When an existing keychain with MD5 cryptographic algorithm has been configured on the RIP interface:

```
GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?: No
Out-of-memory state: Normal
Broadcast for V2: No
Accept Metric 0?: No
Send versions: 2
Receive versions: 2
Interface state: Unknown State
IP address: 0.0.0.0/0
Metric Cost: 0
Split horizon: Enabled
Poison Reverse: Disabled
Socket set options:
Joined multicast group?: No
LPTS filter set?: No
```

```
Authentication mode: MD5 Key chain: <key-chain-name>
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>
```

When the keychain configured on the RIP interface does not exist or does not have any active keys:

```
GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?: No
Out-of-memory state: Normal
Broadcast for V2: No
Accept Metric 0?: No
Send versions: 2
Receive versions: 2
Interface state: Unknown State
IP address: 0.0.0.0/0
Metric Cost: 0
Split horizon: Enabled
Poison Reverse: Disabled
Socket set options:
Joined multicast group?: No
```

```

LPTS filter set?:          No

Authentication mode: MD5 Key chain: <key-chain-name>
No active key found in keychain database.
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>

```

When an active key exists in the keychain configured on the RIP interface, but not configured with MD5 cryptographic algorithm:

```

GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?:             No
Out-of-memory state:     Normal
Broadcast for V2:        No
Accept Metric 0?:       No
Send versions:           2
Receive versions:        2
Interface state:         Unknown State
IP address:              0.0.0.0/0
Metric Cost:             0
Split horizon:           Enabled
Poison Reverse:          Disabled
Socket set options:
Joined multicast group?:  No
LPTS filter set?:        No

Authentication mode: MD5 Key chain: <key-chain-name>
Key(s) not configured with MD5 cryptographic algorithm.
Packets received: <num-rx-packets>
Authenticated packets received: <num-auth-rx-packets>
Packets dropped due to wrong keychain config: <num-rx-wrong-auth-cfg-pkts>
Packets received without authentication data: <num-rx-auth-missing-pkt>
Packets received with invalid authentication: <num-rx-invalid-auth-pkt>

```

When no authentication keychain has been configured on the RIP interface:

```

GigabitEthernet0/3/0/3 (Forward Reference)
Rip enabled?:             No
Out-of-memory state:     Normal
Broadcast for V2:        No
Accept Metric 0?:       No
Send versions:           2
Receive versions:        2
Interface state:         Unknown State
IP address:              0.0.0.0/0
Metric Cost:             0
Split horizon:           Enabled
Poison Reverse:          Disabled
Socket set options:
Joined multicast group?:  No
LPTS filter set?:        No

Authentication mode is not set.
Packets received: <num-rx-packets>

```

This table describes the significant fields shown in the display.

Table 170: show rip [vrf <vrf-name>] interface Field Descriptions

| | |
|------------------------------------|-------------------------------------|
| Authentication mode: MD5 Key chain | MD5 authentication mode is enabled. |
| Current active send key id | Active send key ID. |
| Current active receive key id | Active receive key ID. |

| | |
|--|--|
| Packets received | Number of packets received on the interface. |
| Authenticated packets received | Number packets received with valid authentication. |
| Packets dropped due to wrong keychain config | Number of packets dropped due to wrong keychain configuration. |
| Packets received without authentication data | Number packets received without authentication data . |
| Packets received with invalid authentication | Number of packets received with invalid authentication. |
| No active key found in keychain database | No active keys are available in IOS XR keychain database. |
| Key(s) not configured with MD5 cryptographic algorithm | Keys are not configured with MD5 cryptographic algorithm. |
| Authentication mode is not set | Authentication mode is not set. |

show rip statistics

To display statistical entry information from the Routing Information Protocol (RIP) topology table, use the **show rip statistics** command in EXEC mode.

show rip [**vrf** {*vrf* | **all**}] **statistics**

Syntax Description

| | |
|--|---|
| vrf { <i>vrf</i> all } | (Optional) Specifies a particular VPN routing and forwarding (VRF) instance or all VRF instances. |
|--|---|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|---------|------------|
| rip | read |

Examples

The following example is sample output from the **show rip statistics** command:

```
RP/0/RSP0/CPU0:router# show rip statistics

RIP statistics:
  Total messages sent:          5597
  Message send failures:       0
  Regular updates sent:        5566
  Queries responded to:        0
  RIB updates:                  6
  Total packets received:      5743
  Discarded packets:           0
  Discarded routes:            0
  Number of routes allocated:  18
  Number of paths allocated:   14
```

```

Route malloc failures:    0
Path malloc failures:    0

```

This table describes the significant fields shown in the display.

Table 171: show rip statistics Field Descriptions

| Field | Description |
|----------------------------|---|
| Total messages sent | Number of RIP packets sent. |
| Message send failures | Number of times that the packet send operation failed. |
| Queries responded to | Number of times RIP updates are sent in response to a RIP query. |
| RIB updates | Number of route addition and deletion messages sent to RIB. |
| Total packets received | Number of RIP packets received. |
| Discarded packets | Number of received RIP packets that are discarded. |
| Discarded routes | Number of routes (in received RIP update packets) that are discarded. |
| Number of routes allocated | Number of routes allocated for the RIP internal topology database. |
| Number of paths allocated | Number of paths allocated for the RIP internal topology database. |
| Route malloc failures | Number of failures during route allocation. |
| Path malloc failures | Number of failures during route allocation. |



Note

The number of routes found in the allocated field might be the same number of routes present in the RIP database.

site-of-origin (RIP)

To configure the Site of Origin (SoO) filtering on a Routing Information Protocol (RIP) interface, use the **site-of-origin** command in interface configuration mode. To disable SoO filtering on an interface, use the **no** form of this command.

site-of-origin {*as-number* : *number*| *ip-address* : *number*}

no site-of-origin {*as-number* : *number*| *ip-address* : *number*}

Syntax Description

| | |
|--------------------|---|
| <i>as-number</i> : | Autonomous system number. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. <p>A colon is used to separate the autonomous system number and network number.</p> |
| <i>number</i> | Network number. Range is from 0 to 4294967295 when a 2-byte AS number is used. Range is from 0 to 65535 when a 4-byte AS number is used. |
| <i>ip-address</i> | The IP address argument specifies the IP address in four-part, dotted-decimal notation. A colon is used to separate the IP address and network number. |

Command Default

No default behavior or values

Command Modes

Interface configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A RIP process must be capable of retrieving the SoO attribute on routes redistributed from the Border Gateway Protocol (BGP) when required to support complex topologies that include MPLS VPN links between sites with backdoor links.

Use the **site-of-origin** command to set an SoO BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SoO extended community uniquely identifies the site from which a provider edge (PE) router has learned a route.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to configure SoO filtering on a RIP interface:

```
RP/0/RSP0/CPU0:router(config)# router rip  
RP/0/RSP0/CPU0:router(config-rip) interface GigabitEthernet 0/1/0/0  
RP/0/RSP0/CPU0:router(config-rip-if) site-of-origin 10.0.0.1:20
```

split-horizon disable (RIP)

To disable split horizon for a Routing Information Protocol (RIP) process, use the **split-horizon disable** command in interface configuration mode. To enable split horizon, use the **no** form of this command.

split-horizon disable

no split-horizon disable

Syntax Description This command has no arguments or keywords.

Command Default Split horizon is enabled for a RIP process.

Command Modes Interface configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can explicitly specify the **split-horizon disable** command in your configuration.

If split horizon is disabled, the poison reverse configuration is ignored.



Note In general, we recommend that you do not change the default state of split horizon unless you are certain that your application requires the change to properly advertise routes.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to disable split horizon on a Packet-over-SONET/SDH link:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet 0/1/0/0
RP/0/RSP0/CPU0:router(config-rip-if)# split-horizon disable
```

Related Commands

| Command | Description |
|---|--|
| poison-reverse , on page 1477 | Enables poison reverse processing of RIP router updates. |

timers basic

To adjust Routing Information Protocol (RIP) network timers, use the **timers basic** command in router configuration mode. To restore the timers default values, use the **no** form of this command.

timers basic *update invalid holddown flush*

no timers basic

Syntax Description

| | |
|-----------------|--|
| <i>update</i> | Rate, in seconds, at which updates are sent. This is the fundamental timing parameter of the routing protocol. Range is 5 to 50000. |
| <i>invalid</i> | Interval, in seconds, after which a route is declared invalid; it should be at least three times the value of the update argument. A route becomes invalid when there is an absence of updates that refresh the route. The route then enters into a holddown state. The route is marked inaccessible and is advertised as unreachable. Range is 15 to 200000. |
| <i>holddown</i> | Interval, in seconds, during which routing information regarding better paths is suppressed. It should be at least three times the value of the update argument. A route enters into a holddown state when an update packet is received that indicates that the route is unreachable. The route is marked inaccessible and is advertised as unreachable. When holddown expires, routes advertised by other sources are accepted, and the route is no longer inaccessible. Range is 15 to 200000. |
| <i>flush</i> | Amount of time, in seconds, that must pass before the route is removed from the routing table; the interval specified should be greater than the value of the <i>invalid</i> argument. If it is less than the invalid timer value, the proper holddown interval cannot elapse, which results in a new route being accepted before the holddown interval expires. Range is 16 to 250000. |

Command Default

update : 30
invalid : 180
holddown : 180
flush : 240

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The basic timing parameters for RIP are adjustable. Because RIP is running a distributed, asynchronous routing algorithm, these timers must be the same for all routers in the network.

**Note**

Use the **show rip** command to display the current and default timer values.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to set updates to be broadcast every 5 seconds. If a router is not heard from in 15 seconds, the route is declared unusable. Further information is suppressed for an additional 15 seconds. At the end of the flush period, the route is flushed from the routing table.

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip) timers basic 5 15 15 30
```

Related Commands

| Command | Description |
|---|---|
| show rip , on page 1493 | Displays configuration and status of RIP. |

validate-update-source disable

To stop the Cisco IOS XR software from validating the source IP address of incoming routing updates for Routing Information Protocol (RIP), use the **validate-update-source disable** command in router configuration mode. To reenble this function, use the **no** form of this command.

validate-update-source disable

no validate-update-source disable

Syntax Description This command has no arguments or keywords.

Command Default The source IP address of incoming updates for RIP is always validated.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

When the **validate-update-source disable** command is used, validation is not performed.

By default, the software ensures that the source IP address of incoming routing updates is on the same IP network as one of the addresses defined for the receiving interface.

For unnumbered IP interfaces (interfaces configured as IP unnumbered), no checking is performed.

| Task ID | Task ID | Operations |
|---------|---------|-------------|
| | rip | read, write |

Examples The following example shows how to disable source validation:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip) validate-update-source disable
```

vrf (RIP)

To define a VPN routing and forwarding (VRF) instance and to enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove a VRF instance use the no form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

| | |
|----------|---|
| vrf-name | Specifies a particular VPN routing and forwarding instance. |
|----------|---|

Command Default

No VRFs are defined.

Command Modes

Router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the vrf command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

From VRF configuration mode, you can issue all commands available in router configuration mode such as the **auto-summary** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| rip | read, write |

Examples

The following example shows how to enter VRF configuration mode and identify RIP commands that can be issued from that mode:

```
RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# vrf vpn-1
RP/0/RSP0/CPU0:router(config-rip-vrf)# ?

    auto-summary          Enable automatic network number summarization
```

```
broadcast-for-v2      Send RIP v2 output packets to broadcast address
commit               Commit the configuration changes to running
default-information  Control distribution of default information
default-metric       Set metric of redistributed routes
describe             Describe a command without taking real actions
distance             Define an administrative distance
do                   Run an exec command
exit                 Exit from this submode
interface            Enter the RIP interface configuration submode
maximum-paths        Maximum number of paths allowed per route
neighbor             Specify a neighbor router
no                   Negate a command or set its defaults
nsf                  Enable Cisco Non Stop Forwarding
output-delay         Interpacket delay for RIP updates
redistribute          Redistribute information from another routing protocol
route-policy         Apply route policy to routing updates
show                 Show contents of configuration
timers               Adjust routing timers
validate-update-source Validate source address of routing updates
RP/0/RSP0/CPU0:router(config-rip-vrf)#
```



Routing Policy Language Commands

This module describes the Cisco IOS XR software routing policy language (RPL) commands used to create, modify, monitor, and maintain routing policies.

For detailed information about RPL concepts, configuration tasks, and examples, see the *Implementing Routing Policy on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide* .

- [abort \(RPL\), page 1520](#)
- [add, page 1522](#)
- [apply, page 1524](#)
- [as-path in, page 1526](#)
- [as-path is-local, page 1528](#)
- [as-path length, page 1530](#)
- [as-path neighbor-is, page 1532](#)
- [as-path originates-from, page 1535](#)
- [as-path passes-through, page 1538](#)
- [as-path-set, page 1541](#)
- [as-path unique-length, page 1543](#)
- [community is-empty, page 1545](#)
- [community matches-any, page 1546](#)
- [community matches-every, page 1548](#)
- [community-set, page 1551](#)
- [delete community, page 1554](#)
- [delete extcommunity rt, page 1556](#)
- [destination in, page 1558](#)
- [done, page 1560](#)
- [drop, page 1562](#)

- [edit](#), page 1564
- [end-global](#), page 1567
- [end-policy](#), page 1568
- [end-set](#), page 1570
- [extcommunity rt is-empty](#), page 1572
- [extcommunity rt matches-any](#), page 1574
- [extcommunity rt matches-every](#), page 1576
- [extcommunity rt matches-within](#), page 1578
- [extcommunity-set cost](#), page 1580
- [extcommunity-set rt](#), page 1582
- [extcommunity-set soo](#), page 1584
- [extcommunity soo is-empty](#), page 1586
- [extcommunity soo matches-any](#), page 1588
- [extcommunity soo matches-every](#), page 1590
- [if](#), page 1592
- [local-preference](#), page 1599
- [med](#), page 1601
- [next-hop in](#), page 1603
- [orf prefix in](#), page 1605
- [origin is](#), page 1607
- [pass](#), page 1609
- [path-type is](#), page 1611
- [policy-global](#), page 1613
- [prefix-set](#), page 1615
- [prepend as-path](#), page 1618
- [protocol](#), page 1620
- [rd in](#), page 1622
- [rd-set](#), page 1624
- [replace as-path](#), page 1626
- [rib-has-route](#), page 1628
- [route-has-label](#), page 1630
- [route-policy \(RPL\)](#), page 1631
- [route-type is](#), page 1633

- [rpl editor](#), page 1635
- [rpl maximum](#), page 1637
- [set aigp-metric](#), page 1639
- [set community](#), page 1641
- [set core-tree](#), page 1643
- [set dampening](#), page 1645
- [set eigrp-metric](#), page 1647
- [set extcommunity cost](#), page 1649
- [set extcommunity rt](#), page 1651
- [set ip-precedence](#), page 1653
- [set isis-metric](#), page 1655
- [set label](#), page 1656
- [set label-mode](#), page 1657
- [set level](#), page 1659
- [set local-preference](#), page 1661
- [set med](#), page 1663
- [set metric-type \(IS-IS\)](#), page 1665
- [set metric-type \(OSPF\)](#), page 1667
- [set next-hop](#), page 1669
- [set origin](#), page 1671
- [set ospf-metric](#), page 1673
- [set path-selection](#), page 1674
- [set qos-group \(RPL\)](#), page 1676
- [set rib-metric](#), page 1678
- [set rip-metric](#), page 1679
- [set rip-tag](#), page 1681
- [set rpf-topology](#), page 1682
- [set spf-priority](#), page 1684
- [set tag](#), page 1686
- [set traffic-index](#), page 1688
- [set vpn-distinguisher](#), page 1690
- [set weight](#), page 1692
- [show rpl](#), page 1694

- [show rpl active as-path-set](#), page 1696
- [show rpl active community-set](#), page 1699
- [show rpl active extcommunity-set](#), page 1702
- [show rpl active prefix-set](#), page 1705
- [show rpl active rd-set](#), page 1708
- [show rpl active route-policy](#), page 1710
- [show rpl as-path-set](#), page 1713
- [show rpl as-path-set attachpoints](#), page 1715
- [show rpl as-path-set references](#), page 1718
- [show rpl community-set](#), page 1721
- [show rpl community-set attachpoints](#), page 1723
- [show rpl community-set references](#), page 1726
- [show rpl extcommunity-set](#), page 1729
- [show rpl inactive as-path-set](#), page 1732
- [show rpl inactive community-set](#), page 1735
- [show rpl inactive extcommunity-set](#), page 1738
- [show rpl inactive prefix-set](#), page 1741
- [show rpl inactive rd-set](#), page 1744
- [show rpl inactive route-policy](#), page 1746
- [show rpl maximum](#), page 1749
- [show rpl policy-global references](#), page 1751
- [show rpl prefix-set](#), page 1753
- [show rpl prefix-set attachpoints](#), page 1755
- [show rpl prefix-set references](#), page 1758
- [show rpl rd-set](#), page 1761
- [show rpl rd-set attachpoints](#), page 1763
- [show rpl rd-set references](#), page 1765
- [show rpl route-policy](#), page 1768
- [show rpl route-policy attachpoints](#), page 1771
- [show rpl route-policy inline](#), page 1774
- [show rpl route-policy references](#), page 1776
- [show rpl route-policy uses](#), page 1779
- [show rpl unused as-path-set](#), page 1782

- [show rpl unused community-set](#), page 1785
- [show rpl unused extcommunity-set](#), page 1788
- [show rpl unused prefix-set](#), page 1790
- [show rpl unused rd-set](#), page 1793
- [show rpl unused route-policy](#), page 1795
- [source in](#), page 1798
- [suppress-route](#), page 1800
- [tag](#), page 1802
- [unsuppress-route](#), page 1804
- [vpn-distinguisher is](#), page 1806

abort (RPL)

To discard a route policy or set definition and return to global configuration mode, use the **abort** command in the appropriate configuration mode.

abort

Syntax Description

This command has no keywords or arguments.

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy configuration
Prefix set configuration
Route distinguisher set configuration
AS path set configuration
Community set configuration
Extended community set configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to discard a route policy definition that was started and return to global configuration mode:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# if as-path is-local then
RP/0/RSP0/CPU0:router(config-rpl-if)# abort
RP/0/RSP0/CPU0:router(config)#
```

The following example shows how to discard a prefix set definition that was started and return to global configuration mode:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv4-prefix-examples  
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.1.1,  
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.2.0/24,  
RP/0/RSP0/CPU0:router(config-pfx)# abort  
RP/0/RSP0/CPU0:router(config)#
```

add

To add a value to an Routing Information Protocol (RIP) or Enhanced Interior Gateway Protocol (EIGRP) existing metric, use the **add** command in route-policy configuration mode.

```
add {eigrp-metric bandwidth delay reliability loading max-transmission | rip-metric {number | parameter}}
```

Syntax Description

| | |
|-------------------------|---|
| eigrp-metric | Specifies an EIGRP metric attribute. |
| <i>bandwidth</i> | Bandwidth in kilobits per second. Range is from 0 to 4294967295. |
| <i>delay</i> | Delay in 10-microsecond units. Range is from 0 to 4294967295. |
| <i>reliability</i> | Reliability metric. 255 is 100-percent reliable. Range is from 0 to 255. |
| <i>loading</i> | Effective bandwidth (loading). 255 is 100-percent loaded. Range is from 0 to 255. |
| <i>max-transmission</i> | Maximum transmission of the path. Range is from 0 to 65535. |
| rip-metric | Specifies an RIP metric attribute. |
| <i>number</i> | Value assigned to a four-bit unsigned integer. Range is from 0 to 16. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If the add value is greater than the maximum allowed value, the metric is added. If the resulting metric exceeds the maximum for the routing protocol, then the route is dropped (by the client routing protocol).

Task ID

| Task ID | Operations |
|----------------|-------------------|
| route-policy | read, write |

Examples

The following example shows how to offset the RIP metric value:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# add rip-metric 4
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

The following example shows how to set the EIGRP metric value:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# add eigrp-metric 50000 24000 230 14000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

apply

To execute a parameterized or unparameterized policy from within another policy, use the **apply** command in route-policy configuration mode.

apply *policy_name* [*argument1*, *argument2*, . . . , *argumentN*]

Syntax Description

| | |
|--------------------|--|
| <i>policy_name</i> | Name of a route policy. |
| <i>argument</i> | (Optional) Parameter name. The <i>argument</i> can be a value (for example, '100') or a parameter (for example, '\$parameter') |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.1 | Wildcard support was added for apply policy-names. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **apply** command to execute a policy (either parameterized or unparameterized) from within another policy, which allows for the reuse of common blocks of policy.

Wildcards can be used in apply policy names. This supports the nested wildcard apply scenario. A wildcard is specified by inserting an asterisk (*) in place of one of the portions of the apply policy name; the wildcard indicates that any value for that portion of the apply policy name matches. The nested wildcard apply policy allows wildcard (*) based apply nesting. The wildcard operation permits declaration of a generic apply statement that calls all policies that contain a specific defined set of alphanumeric characters, defined on the router.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the policy CustomerIn applies the route-policy SetLocalPref to conditionally set the local preference on a route. The parameters 20, 30, 40, and 50 are passed to the parameterized policy SetLocalPref, where the local preference is set to:

- 20, if the community 217:20 is present in the route
- 30, if the community 217:30 is present in the route
- 40, if the community 217:40 is present in the route
- 50, if the community 217:50 is present in the route

```
RP/0/RSP0/CPU0:router(config)# route-policy SetLocalPref ($lp0, $lp1, $lp2, $lp3, $lp4)
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any ($lp0:$lp1) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp1
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp2) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp2
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp3) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp3
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any ($lp0:$lp4) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference $lp4
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# route-policy CustomerIn($cust)
RP/0/RSP0/CPU0:router(config-rpl)# apply SetLocalPref ($cust, 20, 30, 40, 50)
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# route-policy Cust_217
RP/0/RSP0/CPU0:router(config-rpl)# apply CustomerIn(217)
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

as-path in

To match the AS path of a route to an AS path set, use the **as-path in** command in route-policy configuration mode.

as-path in {*as-path-set-name*| *inline-as-path-set*| *parameter*}

Syntax Description

| | |
|---------------------------|---|
| <i>as-path-set-name</i> | Name of an AS path set. |
| <i>inline-as-path-set</i> | Inline AS path set. The inline AS path set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path in** command as a conditional expression within an **if** statement to match the AS path of a route to an AS path set. The AS path is a sequence of autonomous system numbers traversed by a route.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **as-path in** command evaluates to true if at least one of the regular expressions defined in the associated AS path set matches the AS path attribute of the route.

In the case where the AS path set is defined but contains no elements in it, the **as-path in** conditional expression command returns false.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

For example, assume we have an AS path set named `my-as-set` defined as follows:

```
RP/0/RSP0/CPU0:router(config)# as-path-set my-as-set
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_12$',
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_13$'
RP/0/RSP0/CPU0:router(config-as)# end-set
```

and the following policy excerpt using an `as-path-set-name` argument:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in my-as-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The `AS path in` condition evaluates to true if one or more of the regular expression matches in the set `my-as-set` match the AS path associated with the route. In the case of a defined but empty AS path set, this operator returns false.

The preceding policy excerpt is equivalent to the following version, which uses an `inline-as-path` set variable:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_12$',ios-regex '_13$') then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path is-local

To determine if this router or another router within this autonomous system or confederation originated a Border Gateway Protocol (BGP) route, use the **as-path is-local** command in route-policy configuration mode.

as-path is-local

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path is-local** command as a conditional expression within an **if** statement to determine if this router (or another router within this autonomous system or confederation) originated the route.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Routes that are locally originated within the autonomous system or confederation carry an empty AS path. For the Border Gateway Protocol (BGP) specification, when a route is advertised across the autonomous system boundary or a confederation boundary, the local autonomous system number or confederation ID is appended to the autonomous system path. The AS path of a locally originated aggregate is also empty unless it has been modified by policy.

The **is-local** operator evaluates to true for autonomous system paths that are empty. An empty AS path is how an AS path that is local to our autonomous system is represented in BGP.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the AS path is local, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path is-local then  
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path length

To compare the number of ASN in the AS path of a Border Gateway Protocol (BGP) route, use the **as-path length** command in route-policy configuration mode.

```
as-path length {eq|is|ge|le} {number|parameter}
```

Syntax Description

| | |
|--------------------------|--|
| eq is ge le | Equal to; greater than or equal to; less than or equal to. |
| <i>number</i> | Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path length** command as a conditional expression within an **if** statement to perform a conditional check based on the length of the AS path.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a specific integer value or a range of integer values specified with the **ge** and **le** operators. Any or all these integers can be parameterized. The operator counts one for each autonomous system in the path. In the case where the route may be aggregated and contain one or more AS sets, the length operator adds one for each set present, the occurrence of an AS set typically indicates that this route is an aggregated route, and the aggregated route had a component route that contained one of the autonomous systems in the set. Likewise, in the case of confederations, a count of one is added for each confederation in the path or each confederation set in the path. A null AS path has a length of zero.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the AS path length equals 10, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path length eq 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

Related Commands

| Command | Description |
|---|---|
| as-path in, on page 1526 | Matches the AS path of a route to an AS path set. |
| as-path originates-from, on page 1535 | Compares an AS path against the AS sequence beginning with the AS number that originated a route |
| as-path passes-through, on page 1538 | Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path. |
| as-path unique-length, on page 1543 | Performs specific checks based on the length of the AS path. |

as-path neighbor-is

To test autonomous system numbers at the head of the AS path against a sequence of one or more values or parameters, use the **as-path neighbor-is** command in route-policy configuration mode.

as-path neighbor-is *as-number-list* [**exact**]

Syntax Description

as-number-list Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

exact (Optional) Specifies that with the **exact** keyword, the *as-number-list* value must identically match the AS path for the route; without the **exact** keyword, any element in the *as-number-list* argument matches one or more occurrences of that element in the AS path for the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path neighbor-is** command as a conditional expression within an **if** statement to test the autonomous system number or numbers at the head of the AS path against a sequence of one or more integral values or parameters. In other words, to test to learn if the sequence of autonomous system numbers matches the path beginning with the neighboring autonomous system from which this route was heard.

**Note**

For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command has an equivalent regular expression (`ios-regex`). For example, AS path neighbor-is '1' would be '^1_'.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path neighbor-is '$asnum' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path neighbor-is '10 20' then
```

These statements evaluate to true when the first autonomous system numbers on the AS path match, in the same order, the supplied parameters or integer values in the **neighbor-is** statement. If the neighboring autonomous system location happens to be an AS-set, the operator evaluates to true if the corresponding argument to the **neighbor-is** operator is an element of the AS-set.

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example,

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10 20' then
```

matches an AS path beginning

```
10 10 10 20 ...
```

and an AS path beginning:

```
10 20 ....
```

With the **exact** keyword, repetitions are not ignored, therefore

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path neighbor-is '10 20' exact then
```

matches the second of these AS paths but not the first.

Related Commands

| Command | Description |
|--|--|
| as-path in , on page 1526 | Matches the AS path of a route to an AS path set. |
| as-path length , on page 1530 | Compares the number of ASN in the AS path of a route. |
| as-path originates-from , on page 1535 | Compares an AS path to the AS sequence beginning with the AS number that originated a route. |

| Command | Description |
|--|---|
| as-path passes-through, on page 1538 | Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path. |
| as-path unique-length, on page 1543 | Performs specific checks based on the length of the AS path. |

as-path originates-from

To compare an AS path against the AS sequence beginning with the AS number that originated a route, use the **as-path originates-from** command in route-policy configuration mode.

as-path originates-from *as-number-list* [**exact**]

Syntax Description

as-number-list Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers.

- Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535.
- Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295.
- Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535.

exact (Optional) Specifies that with the **exact** keyword, the *as-number-list* value must identically match the AS path for the route; without the **exact** keyword, any element in the *as-number-list* argument matches one or more occurrences of that element in the AS path for the route.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path originates-from** command as a conditional expression within an **if** statement to compare an AS path to the autonomous system sequence.

**Note**

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **originates-from** operator is similar to the **neighbor-is** operator, except that it looks at the autonomous system number at the opposite end of the AS path. In other words, it is comparing to the autonomous system that originated the route. It can take numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. When more than one number is specified in the list, the sequence of autonomous system numbers listed must appear as a subsequence in the AS path, with the last number corresponding to the autonomous system that originated the route.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path originates-from '$asnum 11' then
```

The first line of the preceding example evaluates to true if autonomous system 11 originated the route and then advertised it to autonomous system 10, from which the route was eventually propagated to us. In the case where the route has been aggregated, and the location of the originating autonomous system contains an AS-set, the **originates-from** operator evaluates to true if the argument to the **originates-from** operator is contained in the AS-set.

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example,

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' then
```

matches an autonomous system path ending

```
...10 10 10 11
```

and an autonomous system path ending

```
...10 11
```

With the **exact** keyword, repetitions are not ignored, therefore

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path originates-from '10 11' exact then
```

matches the second of these autonomous system paths but not the first.

Related Commands

| Command | Description |
|---|---|
| as-path in , on page 1526 | Matches the AS path of a route to an AS path set. |

| Command | Description |
|--|---|
| as-path length, on page 1530 | Compares the number of ASN in the AS path of a route |
| as-path passes-through, on page 1538 | Verifies if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order |
| as-path unique-length, on page 1543 | Performs specific checks based on the length of the AS path. |

as-path passes-through

To verify if the supplied integer or parameter appears anywhere in the AS path or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path, use the **as-path passes-through** command in route-policy configuration mode.

as-path passes-through *as-number-list* [**exact**]

Syntax Description

| | |
|-----------------------|--|
| <i>as-number-list</i> | Numbers or parameters, enclosed in single quotation marks, that represent a sequence of autonomous system numbers. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| exact | (Optional) Specifies that with the exact keyword, the <i>as-number-list</i> value must identically match the AS path for the route; without the exact keyword, any element in the <i>as-number-list</i> argument matches one or more occurrences of that element in the AS path for the route. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path passes-through** command as a conditional expression within an **if** statement to verify if the specified integer or parameter appears anywhere in the AS path or if the sequence of integers and parameters appears.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **passes-through** operator takes a sequence of integers or parameters, enclosed in single quotation marks, as an argument. It can also take a single integer or parameter as an argument. It evaluates to true if the supplied integer or parameter appears anywhere in the AS path, or if the supplied sequence of integers and parameters appears, in the same order, anywhere in the AS path. This includes the **originates-from** or **neighbor-is** location in the AS path.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following are incomplete configuration examples:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '10' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '$asnum' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '10 11' then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path passes-through '10 $asnum 12' then
```

Without the **exact** keyword, repeated autonomous system numbers in the AS path are ignored. For example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '9 10 11' then
```

matches an AS path containing

```
...9 10 10 10 11 ....
```

and an AS path containing:

```
...9 10 11...
```

With the **exact** keyword, repetitions are not ignored. Therefore:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '9 10 11' exact then
```

matches the second of these AS paths but not the first.

Related Commands

| Command | Description |
|--|--|
| as-path in , on page 1526 | Matches the AS path of a route to an AS path set. |
| as-path length , on page 1530 | Compares the number of ASN in the AS path of a route |
| as-path originates-from , on page 1535 | Compares an AS path to the AS sequence beginning with the AS number that originated a route. |

| Command | Description |
|---|--|
| as-path unique-length, on page 1543 | Performs specific checks based on the length of the AS path. |

as-path-set

To create a named AS path set, use the **as-path-set** command in global configuration mode. To remove the named AS path set, use the **no** form of this command.

as-path-set *name*

no as-path-set *name*

Syntax Description

| | |
|-------------|--------------------------|
| <i>name</i> | Name of the AS path set. |
|-------------|--------------------------|

Command Default

No default behavior or values

Command Modes

global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path-set** command to create a named AS path set.

An AS path set comprises operations for matching an AS path attribute.

This command enters AS path set configuration mode, in which you can use the **ios-regex** keyword to indicate the type of regular expression. Single quotation marks are required around the regular expression.

The inline set form is a parenthesized list of comma-separated expressions.

See the “Understanding Regular Expressions, Special Characters and Patterns” appendix in the *Cisco ASR 9000 Series Aggregation Services Router Getting Started Guide* for information about forming regular expressions.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following is a sample definition of an AS path set named aset1:

```
RP/0/RSP0/CPU0:router(config)# as-path-set aset1  
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_42$',  
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_127$'  
RP/0/RSP0/CPU0:router(config-as)# end-set
```

This AS path set is composed of two elements. When used in a matching operation, this AS path set matches any route whose AS path ends with either the autonomous system number 42 or 127.

The following is a sample of an inline set:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in (ios-regex '_42$', ios-regex$ '_127$')  
RP/0/RSP0/CPU0:router(config-rpl-if)# pass  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

as-path unique-length

To perform specific checks based on the length of the AS path (match against the number of unique ASNs in the AS path), use the **as-path unique-length** command in route-policy configuration mode.

```
as-path unique-length {eq|is|ge|le} {number|parameter}
```

Syntax Description

| | |
|--------------------------|--|
| eq is ge le | Equal to; greater than or equal to; less than or equal to. |
| <i>number</i> | Value assigned to an 11-bit unsigned integer. Range is from 0 to 2047. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **as-path unique-length** command as a conditional expression within an **if** statement to perform a match based on the length of the AS path.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **unique-length** operator is similar to the length operator, except that when an AS path has been padded with the same autonomous system number multiple times, the operator counts only one when the route is padded. Therefore, given an AS path of 333 333 111 222 123 444 444 444, the **unique-length** operator would return a value of 5, whereas the length operator would return a value of 8.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to perform checks based on the AS path length. If the AS path matches the specified values, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path unique-length eq 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length ge 10 then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length le 10 then

RP/0/RSP0/CPU0:router(config-rpl)# if as-path unique-length eq $integerparam then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length ge $geparam then
RP/0/RSP0/CPU0:router(config-rpl-if)# if as-path unique-length le $leparam then

RP/0/RSP0/CPU0:router(config-rpl)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl)# endif
```

Related Commands

| Command | Description |
|--|---|
| as-path length, on page 1530 | Performs conditional checks based on the length of the AS path. |

community is-empty

To check if a route has no community attributes associated with it, use the **community is-empty** command in route-policy configuration mode.

community is-empty

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community is-empty** command as a conditional expression within an **if** statement to check if a route has community attributes associated with it.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes no arguments and evaluates to true only if the route has no community attributes associated with it.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, if the route has no community attributes associated with it, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if community is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

community matches-any

To match any elements of a community set, use the **community matches-any** command in route-policy configuration mode.

community matches-any {*community-set-name*| *inline-community-set*| *parameter*}

Syntax Description

| | |
|-----------------------------|---|
| <i>community-set-name</i> | Name of a community set. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community matches-any** command as a conditional expression within an **if** statement to match any element of a community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one community element of the community attribute for the route matches an element in the community set operand. If no community in the route matches any of the specifications in the named or inline set, then the condition evaluates to false. Likewise, when there is no community at all in the route, the condition evaluates to false.

Matching of a community in the route to a specification in a named or an inline set is intuitive. If the community specification in a set is the familiar colon-separated decimal 16-bit numbers specification, or one of the well-known communities, the community matches the specification if the specification denotes the same 32-bit number as that in the route. If the community specification uses a wildcard, then the community in the route matches if it is one of the many communities denoted by the wildcard specification. In inline sets,

community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Communities may also be matched using range and regular expression operators. Range specifications are entered as follows: [*low-value* .. *high-value*]. Either or both colon-separated halves of a community value may contain a range. The following are valid range specifications:

```
10:[100..1000]
[10..100]:80
[10..100]:[100..2000]
```

In addition, the **private-as** keyword may be used to specify the range from 64512 to 65534. Regular expressions are specified as the **ios-regex** keyword followed by a valid regular expression string.

Community values from the route are matched one at a time to the match specifications. Therefore, regex match specifications are expected to represent one individual community value and not a sequence of community values.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, a named community set called my-community-set and a route policy called community-matches-any-example are created. The policy sets the local-preference to 100 for any route that has one or more of the communities in the my-community-set community set. If the route does not have any of these communities, the policy checks whether it has any communities whose first half is in the range from 10 to 25 and whose second half is the value 35, in which case it sets the local-preference to 200. Otherwise, it checks for a community value in the range of 30:100 to 30:500, in which case it sets the local-preference to 300.

```
RP/0/RSP0/CPU0:router(config)# community-set my-community-set
RP/0/RSP0/CPU0:router(config-comm)# 10:20,
RP/0/RSP0/CPU0:router(config-comm)# 10:30,
RP/0/RSP0/CPU0:router(config-comm)# 10:40
RP/0/RSP0/CPU0:router(config-comm)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy community-matches-any-example
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any my-community-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif community matches-any ([10..25]:35) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-any (30:[100..500])
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|---|
| community matches-every , on page 1548 | Matches every element of a community set. |

community matches-every

To match every element of a community set, use the **community matches-every** command in route-policy configuration mode.

community matches-every {*community-set-name*| *inline-community-set*| *parameter*}

Syntax Description

| | |
|-----------------------------|---|
| <i>community-set-name</i> | Name of a community set. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **community matches-every** command as a conditional expression within an **if** statement to match every element of a community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every specification in the named set or inline set specified matches at least one community value in the route. If any community specification in the named or inline set is not matched, then the operation evaluates to false.

Matching of a community in the route to a specification in a named or an inline set is intuitive. If the community-specification in a set is the familiar colon-separated decimal 16-bit numbers specification, or one of the well-known communities, the community matches the specification if the specification denotes the same 32-bit number as that in the route. If the community specification uses a wildcard, then the community in the route matches if it is one of the many communities denoted by the wildcard specification. In inline sets,

community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Communities may also be matched using range and regular expression operators. Range specifications are entered as follows: [*low-value* .. *high-value*]. Either or both colon-separated halves of a community value may contain a range. The following are valid range specifications:

```
10:[100..1000]
[10..100]:80
[10..100]:[100..2000]
```

Therefore, a **matches-every** operation with two community range specifications means that a community must be present in the route that corresponds to each range. For example, in the following statement:

```
if community matches-every (10:[100..200],20:[100..200]) then
```

the statement evaluates as true if one or more communities in the route lie in the range 10:[100..200] and one or more communities in the route lie in the range 20:[100..200].

In addition, the **private-as** keyword may be used to specify the range from 64512 to 65534.

Regular expressions are specified as the **ios-regex** keyword followed by a valid single-quoted regular expression string. Community values from the route are matched one at a time against the match specifications. Therefore, regex match specifications are expected to represent one individual community value and not a sequence of community values.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the route policy named community-matches-every-example sets the local-preference value to 100 for all routes that have all three communities in the my-community-set community set. Routes that do not have all three communities but have a community that matches the first regular expression match have the local-preference value set to 200. Finally, any remaining routes that match the last regular expression have the local-preference values set to 300.

```
RP/0/RSP0/CPU0:router(config)# community-set my-community-set
RP/0/RSP0/CPU0:router(config-comm)# 10:20,
RP/0/RSP0/CPU0:router(config-comm)# 10:30,
RP/0/RSP0/CPU0:router(config-comm)# 10:40
RP/0/RSP0/CPU0:router(config-comm)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy community-matches-every-example
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-every my-community-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-every (ios-regex
'10:[0-9]0_') then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif community matches-every
(ios-regex'20:[0-9]0_') then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|---|---|
| community matches-any, on page 1546 | Matches any element of a community set. |

community-set

To define a community set, use the **community-set** command in global configuration mode. To remove the community set, use the **no** form of this command.

community-set *name*

no community-set *name*

Syntax Description

| | |
|-------------|----------------------------|
| <i>name</i> | Name of the community set. |
|-------------|----------------------------|

Command Default

No default behavior or values

Command Modes

global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Regular expressions and ranges can be specified to match the communities. An attempt to use a community set that contains a range or regular expression to set a community value is rejected when an attempt to attach such a policy is made.

A community set holds community values for matching against the Border Gateway Protocol (BGP) community attribute. A community is a 32-bit quantity. For notational convenience, each community value must be split in half and expressed as two unsigned decimal integers in the range from 0 to 65535, separated by a colon.

The inline form of a community set also supports parameterization. Each 16-bit portion of the community may be parameterized.

The routing policy language (RPL) provides symbolic names for the standard well-known community values: **accept-own** is 0xFFFF0001, **internet** is 0:0, **no-export** is 65535:65281, **no-advertise** is 65535:65282, and **local-as** is 65535:65283.

RPL also provides a facility for using wildcards in community specifications. A wildcard is specified by inserting an asterisk (*) in place of one of the 16-bit portions of the community specification, which indicates that any value for that portion of the community matches.

Every community set must contain at least one community value. An empty community set is invalid and the policy configuration system rejects it.

Community sets can be entered in these formats:

| Format | Description |
|---------------------|--|
| <i>#-remark</i> | Remark beginning with '#' |
| * | Wildcard (any community or part thereof) |
| <i>0-65535</i> | 16-bit half-community number |
| [| Left bracket to begin range |
| accept-own | Accept-Own (BGP well-known community) |
| dfa-regex | DFA (deterministic finite automata) style regular expression |
| internet | Internet (BGP well-known community) |
| ios-regex | Traditional IOS style regular expression |
| local-AS | Do not send outside local AS (BGP well-known community) |
| no-advertise | Do not advertise to any peer (BGP well-known community) |
| no-export | Do not export to next AS (BGP well-known community) |
| private-as | Match within BGP private AS range [64512..65534] |

**Note**

The dfa-regex and ios-regex syntax for community set is "[^:;<>]*:[^:;<>]*[']". This means that regex starts with a single-quote (") followed by a string of any character (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by a colon, and a string of any characters (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by single-quote.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, a community set named cset_accept_own is created:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#community-set cset_accept_own
RP/0/RSP0/CPU0:router(config-comm)#accept-own
RP/0/RSP0/CPU0:router(config-comm)#end-set
```

In the following example, a community set named cset1 is created:

```
RP/0/RSP0/CPU0:router(config)# community-set cset1
```

```
RP/0/RSP0/CPU0:router(config-comm)# 12:34,  
RP/0/RSP0/CPU0:router(config-comm)# 12:56,  
RP/0/RSP0/CPU0:router(config-comm)# 12:78,  
RP/0/RSP0/CPU0:router(config-comm)# internet  
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

In the following example, a community set named cset2 is created:

```
RP/0/RSP0/CPU0:router(config)# community-set cset2  
RP/0/RSP0/CPU0:router(config-comm)# 123:456,  
RP/0/RSP0/CPU0:router(config-comm)# no-advertise,  
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

In the following example, a community set named cset3 is created. This policy uses wildcards and matches all communities where the autonomous system part of the community is 123.

```
RP/0/RSP0/CPU0:router(config)# community-set cset3  
RP/0/RSP0/CPU0:router(config-comm)# 123:*  
RP/0/RSP0/CPU0:router(config-comm)# end-set
```

delete community

To delete community attributes associated with a Border Gateway Protocol (BGP) route, use the **delete community** command in route-policy configuration mode.

```
delete community {all| in {community-set-name| inline-community-set| parameter}| not in
{community-set-name| inline-community-set| parameter}}
```

Syntax Description

| | |
|-----------------------------|--|
| all | Removes all communities except the well-known communities. |
| in | Removes any communities associated with the route that are listed in either the named community set or the inline community set. |
| <i>community-set-name</i> | Name of a community set. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| not in | Removes all communities that are not listed in either the named community set or the inline community set, and are not well-known communities. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **delete community** command to delete community attributes associated with a BGP route.



Note

The **delete community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Communities are 32-bit values carried in Border Gateway Protocol (BGP) routes. Each route may have zero or more communities in an unordered list.

You can remove a well-known community (internet, no-export, no-advertise, or local-as) from a route, but this removal must be done explicitly. This command should be used with a degree of caution. In general, few circumstances exist in which you would need to remove a well-known community.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to delete any communities associated with the routes that are listed in either the named community set or inline community set, respectively.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in my_community_set
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (10:[0..50],20:[60..80])
```

The following example shows how to remove all communities including well-known communities.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (internet, no-export, no-advertise,
local-as, *:*)
```

The following example shows how to remove all communities except for the well-known communities.

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community all
```

The following example shows how to delete the well-known community value internet from a route:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete community in (internet)
```

delete extcommunity rt

To delete route target (RT) extended community attributes associated with a Border Gateway Protocol (route), use the **delete extcommunity rt** command in route-policy configuration mode.

```
delete extcommunity rt {all| in {extcommunity-set-name| inline-extcommunity-set| parameter}| not in {extcommunity-set-name| inline-extcommunity-set| parameter}}
```

Syntax Description

| | |
|--------------------------------|--|
| all | Removes all extended communities. |
| in | Removes any extended communities associated with the routes that are listed in either the named extended community set or the inline extended community set. |
| <i>extcommunity-set-name</i> | Name of an extended community set. |
| <i>inline-extcommunity-set</i> | Inline extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| not in | Removes all extended communities that are not listed in either the named extended community set or the inline extended community set, and are not well-known extended communities. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **delete extcommunity rt** command to delete extended community values from a BGP route target extended community list in a route.



Note The **delete extcommunity rt** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Extended communities are similar to regular Border Gateway Protocol (BGP) communities but contain more data and have a richer structure for encoding information in them.

Extended communities can be in the following forms: SoO:AS:tag, SoO:IP:tag, RT:AS:tag, or RT:IP:tag.

Wildcards (*) and regular expressions are allowed for extended community set elements.

The forms of this command that take a named extended community set or an inline extended community set value as arguments are equivalent. They delete any extended communities that are listed in either the named set or the inline set, respectively.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, all extended communities are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt all
```

In this example, any extended communities that are listed in my-extcommunity-set are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt in my-extcommunity-set
```

In this example, extended communities associated with the route listed in the named inline extended community sets are deleted:

```
RP/0/RSP0/CPU0:router(config-rpl)# delete extcommunity rt in (67:29, 67:55)
```

destination in

To match a destination entry in a named prefix set or inline prefix set, use the **destination in** command in route-policy configuration mode.

destination in {*prefix-set-name*| *inline-prefix-set*| *parameter*}

Syntax Description

| | |
|-------------------------------|---|
| <i>prefix-set-name</i> | Name of a prefix set. |
| <i>inline-prefix-set</i> | Inline prefix set. The inline prefix set must be enclosed in parentheses. |
| <i>parameter</i> parameter | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **destination in** command as a conditional expression within an **if** statement to match a destination entry in a named prefix set or inline prefix set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a named prefix set or an inline prefix set value as an argument. The condition returns true if the destination entry matches any entry in the prefix set or inline prefix set. An attempt to match a destination using a prefix set that is defined but contains no elements returns false.

The routing policy language (RPL) provides the ability to test destinations for a match to a list of prefix match specifications using the **in** operator. The **destination in** command is protocol-independent.

In Border Gateway Protocol (BGP), the destination of a route is also known as its network-layer reachability information (NLRI). It comprises a prefix value and a mask length.

RPL supports both 32-bit IPv4 prefixes, specified in dotted-decimal format, and 128-bit IPv6 prefixes, specified in colon-separated hexadecimal format.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, a prefix set named my-prefix-set is defined and a route policy named use-destination-in is created. Within the use-destination-in route policy, the **destination in** command is used within an **if** statement to learn if the destination is in the prefix-set named my-prefix-set. If it is, then local preference is set to 100. If it is not in my-prefix-set but does match the next prefix specifications, then local preference is set to 200.

```
RP/0/RSP0/CPU0:router(config)# prefix-set my-prefix-set
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.0.1/32,
RP/0/RSP0/CPU0:router(config-pfx)# fe80::203:0:0:0/64,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.0.2/24 le 32
RP/0/RSP0/CPU0:router(config-pfx)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy use-destination-in
RP/0/RSP0/CPU0:router(config-rpl)# if destination in my-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif destination in (10.0.0.1/32, 10.0.0.2/24 le
32) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

In the following example, a prefix set named ipv6-prefix-set is defined and a route policy named ipv6-destination-in is created. Within the ipv6-destination-in route policy, the **destination in** command is used within an **if** statement to learn if the destination is in the prefix-set named ipv6-prefix-set. If it is, then the next-hop is set to 2001:abcd:fedc::1. If it is not in ipv6-prefix-set but does match the next prefix specifications, then the next-hop is set to 1111:2222:3333:4444:5555:6666:7777:8888.

```
RP/0/RSP0/CPU0:router(config)# prefix-set ipv6-prefix-set
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:1::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:2::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:3::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:4::/64
RP/0/RSP0/CPU0:router(config-pfx)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy ipv6-destination-in
RP/0/RSP0/CPU0:router(config-rpl)# if destination in ipv6-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set next-hop 2001:abcd:fedc::1
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif destination in (2001::1, 2002:1:2:3::/64)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set next-hop
1111:2222:3333:4444:5555:6666:7777:8888
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

done

To stop executing a policy and accept the route, use the **done** command in route-policy configuration mode.

done

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **done** command to stop executing the policy and accept the route.



Note

The **done** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

When encountering a **done** statement the route is passed and no further policy statements are executed. All modifications made to the route prior to the **done** statement are still valid.



Note

The default action of a route policy is to drop or discard any routes that have not been either explicitly passed or for which no attempt has been made to modify with an action. The routing policy language (RPL) does not have specific “match clauses,” which means the default drop behavior is controlled by whether a route has been explicitly passed or an attempt has been to modify the route using an action statement.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the destination match succeeds for 29.0.0.0/8 le 32, the execution continues past set community 102:12 and onto the next statement. If the destination match succeeds for 39.0.0.0/8 le 32 execution, then the policy execution stops when it encounters the *done* statement.

```
RP/0/RSP0/CPU0:router(config)# route-policy done_st_example
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (29.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:12
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (39.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:39
RP/0/RSP0/CPU0:router(config-rpl-if)# done
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (49.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:49
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (59.0.0.0/8 le 32) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community 102:59
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

drop

To discard a route, use the **drop** command in route-policy configuration mode.

drop

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **drop** command within a route policy to drop a route.



Note

The **drop** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command causes the route to be dropped. After a route is dropped, no further execution of policy occurs. Therefore, if after executing the first two statements of a policy the **drop** statement is encountered, the route is discarded and execution stops immediately even when the policy contains further statements.



Note

The default action of a route policy is to drop or discard any routes that have not been either explicitly passed or attempted to be modified with an action. The routing policy language (RPL) does not have specific “match clauses,” which means the default drop behavior is controlled by whether a route has been explicitly passed or an attempt has been to modify the route using an action statement.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, any route with a destination address contained within the prefix set pset1 is dropped:

```
RP/0/RSP0/CPU0:router(config-rpl)# if destination in pset1 then  
RP/0/RSP0/CPU0:router(config-rpl-if)# drop  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

edit

To edit the contents of a route policy, a prefix set, an AS path set, a community set, or an extended community set, use the **edit** command in EXEC configuration mode.

```
edit {route-policy|prefix-set|as-path-set|community-set|extcommunity-set {rt|soo}|policy-global|rd-set} name [nano|emacs|vim|inline {add|prepend|remove} set-element]
```

Syntax Description

| | |
|-------------------------|---|
| route-policy | Edits the contents of a route policy. |
| prefix-set | Edits the contents of a prefix set. |
| as-path-set | Edits the contents of an AS path set. |
| community-set | Edits the contents of a community set. |
| extcommunity-set | Edits the contents of an extended community set of the specified type. |
| rt | Edits the BGP route target (RT) extended community. |
| soo | Edits the BGP site of origin (SoS) extended community. |
| policy-global | Edits the contents of policy-global definitions. |
| rd-set | Edits the contents of a route-distinguisher set. |
| <i>name</i> | Name of a route policy, a prefix set, an AS path set, a community set, or an extended community set, RD set, or global parameters. |
| nano | (Optional) Uses GNU Nano text editor. |
| emacs | (Optional) Uses Micro Emacs editor. |
| vim | (Optional) Uses VI Improved editor. |
| inline | (Optional) Uses the command line. |
| add | Appends the element to the set. |
| prepend | Prepends the element to the set. |
| remove | Removes the element from the set. |
| <i>set-element</i> | Value of the set element. Note To inline edit multiple set elements separated with comma, use quotes to club the entries as a single argument. Example: edit extcommunity-set rt rt_set inline add "4:4,5:4" |

Command Default Default editor is GNU nano text editor

Command Modes EXEC configuration

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **edit** command to edit the contents of a route policy, a prefix set, an AS path set, a community set, an extended community set, a global policy, or a route destination set.

After editing with Nano, save the edit buffer and exit the editor using the Ctrl-X keystroke.

After editing with Emacs, save the editor buffer by using the Ctrl-X and Ctrl-S keystrokes. To save and exit the editor, use the Ctrl-X and Ctrl-C keystrokes.

After editing with VIM, to write to a current file and exit use the :wq or :x or ZZ keystrokes. To quit and confirm, use the :q keystrokes. To quit and discard changes, use the :q! keystrokes.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read, write |

Examples In the following example, the policy_A policy is opened in the editor:

```
RP/0/RSP0/CPU0:router# edit route-policy policy_A
-----
== MicroEMACS 3.8b () == rpl_edit.139281 ==
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

== MicroEMACS 3.8b () == rpl_edit.139281 ==
Parsing.
83 bytes parsed in 1 sec (82)bytes/sec
Committing.
1 items committed in 1 sec (0)items/sec
Updating.
Updated Commit database in 1 sec
```

If there are parse errors, you are asked whether editing should continue:

```
RP/0/RSP0/CPU0:router#edit route-policy policy_B
```

```
== MicroEMACS 3.8b () == rpl_edit.141738
route-policy policy_B
  set metric-type type_1
  if destination in (2001::/8) then
    drop
  endif
end-policy
!
== MicroEMACS 3.8b () == rpl_edit.141738 ==
Parsing.
105 bytes parsed in 1 sec (103)bytes/sec

% Syntax/Authorization errors in one or more commands.!! CONFIGURATION
FAILED DUE TO SYNTAX/AUTHORIZATION ERRORS
  set metric-type type_1
  if destination in (2001::/8) then
    drop
  endif
end-policy
!

Continue editing? [no]:
```

If you answer **yes**, the editor continues on the text buffer from where you left off. If you answer **no**, the running configuration is not changed and the editing session is ended.

After the policy is opened, it may be manipulated using normal editor commands, then saved and committed to the running configuration.

end-global

To end the definition of global parameters and exit global parameter configuration mode, use the **end-global** command in global parameter configuration mode.

end-global

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Global parameter configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-global** command to end the definition of global parameters and exit global parameter configuration mode.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, the **end-global** command ends the definition of global parameters:

```
RP/0/RSP0/CPU0:router(config)#policy-global
RP/0/RSP0/CPU0:router(config-rp-gl)# glbpathype 'ebgp'
RP/0/RSP0/CPU0:router(config-rp-gl)# glbtags '100'
RP/0/RSP0/CPU0:router(config-rp-gl)# end-global
```

| Related Commands | Command | Description |
|------------------|---|---|
| | policy-global, on page 1613 | Enters global parameter configuration mode. |

end-policy

To end the definition of a route policy and exit route-policy configuration mode, use the **end-policy** command in route-policy configuration mode.

end-policy

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-policy** command to end the definition of a route policy and exit route-policy configuration mode.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the **end-policy** command ends the definition of a route policy:

```
RP/0/RSP0/CPU0:router(config)#route-policy med-to-local-pref
RP/0/RSP0/CPU0:router(config-rpl)#if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 10
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif med eq 200 then
RP/0/RSP0/CPU0:router(config-elseif)# set local-preference 60
RP/0/RSP0/CPU0:router(config-elseif)# elseif med eq 250 then
RP/0/RSP0/CPU0:router(config-elseif)# set local-preference 0

RP/0/RSP0/CPU0:router(config-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|---|--|
| route-policy (RPL) , on page 1631 | Defines a route policy and enters route-policy configuration mode. |

end-set

To end the definition of an AS path set, a prefix set, a community set, an extended community set, or an RD set and return to global configuration mode, use the **end-set** command in route-policy configuration mode.

end-set

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values

Command Modes

AS path set configuration
Prefix set configuration
Community set configuration
Extended community set configuration
Route distinguisher set configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **end-set** command to end the definition of an AS path set, a prefix set, a community set, or an extended community set.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the **end-set** command ends the definition of an AS path set named aset1:

```
RP/0/RSP0/CPU0:router(config)# as-path-set aset1
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_42$',
RP/0/RSP0/CPU0:router(config-as)# ios-regex '_127$'

RP/0/RSP0/CPU0:router(config-as)# end-set
```

```
RP/0/RSP0/CPU0:router(config)#
```

The following example shows how to create an RD set called `my_rd_set` and use the `end-set` command to end the definition:

```
RP/0/RSP0/CPU0:router(config)# rd-set my_rd_set  
RP/0/RSP0/CPU0:router(config-rd)# 172.16.0.0/16:*,  
RP/0/RSP0/CPU0:router(config-rd)# 172.17.0.0/16:100,  
RP/0/RSP0/CPU0:router(config-rd)# 192:*,  
RP/0/RSP0/CPU0:router(config-rd)# 192:100  
RP/0/RSP0/CPU0:router(config-rd)# end-set
```

extcommunity rt is-empty

To check if a Border Gateway Protocol (BGP) route has route target (RT) extended community attributes associated with it, use the **extcommunity rt is-empty** command in route-policy configuration mode.

extcommunity rt is-empty

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or value

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt is-empty** command as a conditional expression within an **if** statement to check if a BGP route has extended community attributes associated with it.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **is-empty** operator takes no arguments and evaluates to true if the route has no extended community attributes associated with it.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the extended community is empty, then the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

extcommunity rt matches-any

To match any element of a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity rt matches-any** command in route-policy configuration mode.

extcommunity rt matches-any {*extcommunity-set-name*|*inline-extcommunity-set*|*parameter*}

Syntax Description

| | |
|--------------------------------|--|
| <i>extcommunity-set-name</i> | Name of an RT extended community set. |
| <i>inline-extcommunity-set</i> | Inline RT extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt matches-any** command as a conditional expression within an **if** statement to match elements of an extended community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one extended community in the route matches an extended community specification in the named or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, an extended community set named my-extcommunity-set and a parameterized route-policy named my-extcommunity-set-example(\$tag,\$ip) are defined. The **extcommunity rt matches-any** command is used in an if statement such that if at least one extended community in the route matches an extended community specification in the named set, then the local preference is set to 100. If there is no extended community in the route that matches any of the specifications in the named set, then the condition evaluates as false and the extended community is compared to the inline extended sets.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RSP0/CPU0:router(config-ext)# 10:615,
RP/0/RSP0/CPU0:router(config-ext)# 10:6150,
RP/0/RSP0/CPU0:router(config-ext)# 15.15.15.15:15
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt matches-any my-extcommunity-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-any (10:20, 10:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any ($ip:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity rt matches-any (2.3.4.5:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|---|---|
| extcommunity rt matches-every, on page 1576 | Matches every element of a BGP RT extended community set. |

extcommunity rt matches-every

To match every element of a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity rt matches-every** command in route-policy configuration mode.

extcommunity rt matches-every {*extcommunity-set-name*| *inline-extcommunity-set*| *parameter*}

Syntax Description

| | |
|--------------------------------|--|
| <i>extcommunity-set-name</i> | Name of an RT extended community set. |
| <i>inline-extcommunity-set</i> | Inline RT extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt matches-every** command as a conditional expression within an **if** statement to match every element of an RT extended community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every extended community value in the extended community attribute for the route matches at least one element of the extended community set or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, an extended community set named my-extcommunity-set and a parameterized route policy named extcommunity-matches-every-example (\$as, \$tag) are defined. The condition extcommunity rt matches-every is used in an if statement in this policy. If it evaluates to true, the local-preference value is set to 100. If it evaluates to false, the extended community is evaluated using an inline set. If that condition evaluates to true, the local-preference value is set to 200. If it evaluates to false, the local-preference value is set to 300.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt my-extcommunity-set
RP/0/RSP0/CPU0:router(config-ext)# 10:20,
RP/0/RSP0/CPU0:router(config-ext)# 10:30,
RP/0/RSP0/CPU0:router(config-ext)# 10:40
RP/0/RSP0/CPU0:router(config-ext)# end-set
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as,$tag)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity rt matches-every my-extcommunity-set
then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity rt matches-every (10:20, 10:$tag,
  $as:30) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|---|---|
| extcommunity rt matches-any, on page 1574 | Matches any element of a BGP RT extended community set. |

extcommunity rt matches-within

To match at least one element of an extended community set of a Border Gateway Protocol (BGP) route target (RT), use the **extcommunity rt matches-within** command in route-policy configuration mode.

```
extcommunity rt matches-within {rt-type-extcommunity-set-name| inline-extcommunity-set| parameter}
```

Syntax Description

| | |
|--------------------------------------|--|
| <i>rt-type-extcommunity-set-name</i> | Name of an RT extended community set. |
| <i>inline-extcommunity-set</i> | Inline RT extended community set, enclosed in parentheses. |
| <i>parameter</i> | Parameter name preceded with a "\$" symbol. |

Command Default

None

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity rt matches-within** command as a conditional expression within an if statement to match elements of an extended community set.



Note

For a list of all conditional expressions available within an if statement, see the if command.

A simple condition using the matches-within operator evaluates as true if all the elements in extended community from the route match any element in the extended community set. For example, let 'c' be the RTs from the route and 'm' be the RT set from the policy. With the **extcommunity rt matches-within** configuration, each value in 'c' must match any (or at least one) value in 'm'.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

| Task ID | Task ID | Operation |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples

In the following example, an extended community set named *my-extcommunity-set* and a parameterized route-policy named *my-extcommunity-set-example(\$tag,\$ip)* are defined. The **extcommunity rt matches-within** command is used in an if statement such that if all extended community values in the route match any element of the extended community specification in the named set, then the local preference is set to 100.

```
RP/0/RSP0/CPU0:router(config)#extcommunity-set rt my-extcommunity-set
RP/0/RSP0/CPU0:router(config-ext)#10:615,
RP/0/RSP0/CPU0:router(config-ext)#10:6150,
RP/0/RSP0/CPU0:router(config-ext)#15.15.15.15:15
RP/0/RSP0/CPU0:router(config-ext)#end-set
RP/0/RSP0/CPU0:router(config)#route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RSP0/CPU0:router(config-rpl)#if extcommunity rt matches-within my-extcommunity-set
then
RP/0/RSP0/CPU0:router(config-rpl-if)#set local-preference 100
```

extcommunity-set cost

To define a cost extended community set, use the **extcommunity-set cost** command in global configuration mode. To remove the cost extended community set, use the **no** form of this command.

extcommunity-set cost *name*

no extcommunity-set cost *name*

Syntax Description

| | |
|-------------|---|
| <i>name</i> | Name of a cost extended community set. The <i>name</i> argument is case sensitive, can contain any alphanumeric characters, and can be up to 63 characters in length. |
|-------------|---|

Command Default

No default behavior or values

Command Modes

global configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support was added for more cost extended community formats. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity-set cost** command to define a cost extended community set.

An extended community set is analogous to a community set except that it contains extended community values instead of regular community values. Extended community values are 64-bit structured values. An extended community set also supports named forms and inline forms.

Cost extended communities can be entered in these formats:

- **#-remark** ---Remark beginning with '#'
- **0-255**---Decimal number
- **abort** ---Discard RPL definition and return to top level config
- **end-set** ---End of set definition
- **exit** ---Exit from the submode
- **igp**:---Cost Community with IGP as point of insertion
- **pre-bestpath**: ---Cost Community with Pre-Bestpath as point of insertion

- **show** ---Show partial RPL configuration

Multiple cost community set clauses can be configured in each route policy block or sequence. Each cost community set clause must have a different ID (0-255). The cost community set clause with the lowest cost-value is preferred by the best path selection process when all other attributes are equal.

As with community sets, the inline form supports parameterization within parameterized policies. Either portion of the extended community value can be parameterized.

Every extended community set must contain at least one extended community value. Empty extended community sets are invalid and the policy configuration system rejects them.

Wildcards (*) and regular expressions are allowed for extended community set elements.

Examples

In the following example, a cost extended community set named extcomm-cost is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set cost extcomm-cost  
RP/0/RSP0/CPU0:router(config-ext)# IGP:90:914,  
RP/0/RSP0/CPU0:router(config-ext)# Pre-Bestpath:91:915  
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set rt

To define a Border Gateway Protocol (BGP) route target (RT) extended community set, use the **extcommunity-set rt** command in global configuration mode. To remove the RT community set, use the **no** form of this command.

extcommunity-set rt *name*

no extcommunity-set rt *name*

Syntax Description

| | |
|-------------|---------------------------------------|
| <i>name</i> | Name of an RT extended community set. |
|-------------|---------------------------------------|

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support was added for more rt extended community formats. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity-set rt** command to define an RT extended community set for BGP.

Regular expressions and ranges can be specified to match the extended communities. Regular expressions and ranges can be specified in an extended community set to support the matching of communities. An attempt to use an extended community set that contains a range or regular expression to set an extended community set value is rejected when an attempt to attach such a policy is made.

An extcommunity set RT holds RT extended community values to match against the Border Gateway Protocol (BGP) RT extended community attribute. RT extended communities can be entered in these formats:

- *#-remark* ---Remark beginning with '#'
- **---* Wildcard (any community or part thereof)
- *1-4294967295* ---32-bit decimal number
- *1-65535* ---16-bit decimal number
- *A.B.C.D/M:N* ---Extended community - IPv4 prefix format

- *A.B.C.D:N*---Extended community - IPv4 format
- *ASN:N* ---Extended community - ASPLAIN format
- *X.Y:N* ---Extended community - ASDOT format
- **dfa-regex** ---DFA (deterministic finite automata) style regular expression
- **ios-regex** ---Traditional IOS style regular expression

**Note**

The dfa-regex and ios-regex syntax for community set is "[^:;<>]*:[^:;<>]*[]". This means that regex starts with a single-quote (") followed by a string of any character (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by a colon, and a string of any characters (that does not include single-quote, colon, ampersand, less-than, greater-than, or space) followed by single-quote.

N is a number within the range of 1 to 65535.

Examples

In the following example, an RT extended community set named extcomm-rt is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set rt extcomm-rt  
RP/0/RSP0/CPU0:router(config-ext)# 10002:666  
RP/0/RSP0/CPU0:router(config-ext)# 10.0.0.2:666  
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity-set soo

To define a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity-set soo** command in global configuration mode. To remove the SoO extended community set, use the **no** form of this command.

extcommunity-set soo *name*

no extcommunity-set soo *name*

Syntax Description

| | |
|-------------|--|
| <i>name</i> | Name of an SoO extended community set. |
|-------------|--|

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|--|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Support was added for more soo extended community formats. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity-set soo** command to define an SoO extended community set.

An extcommunity set soo holds SoO extended community values to match against the Border Gateway Protocol (BGP) SoO extended community attribute. SoO extended communities can be entered in these formats:

- *#-remark* ---Remark beginning with '#'
- **---* Wildcard (any community or part thereof)
- *1-4294967295---*32-bit decimal number
- *1-65535* ---16-bit decimal number
- *A.B.C.D/M:N* ---Extended community - IPv4 prefix format
- *A.B.C.D:N*---Extended community - IPv4 format
- *ASN:N* ---Extended community - ASPLAIN format
- *X.Y:N* ---Extended community - ASDOT format

- **abort** ---Discard RPL definition and return to top level config
- **dfa-regex** ---DFA style regular expression
- **end-set** ---End of set definition
- **exit** ---Exit from the submode
- **ios-regex** ---Traditional IOS style regular expression
- **show** ---Show partial RPL configuration

N is a site-specific number.

Examples

In the following example, a SoO extended community set named extcomm-soo is defined:

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo extcomm-soo
RP/0/RSP0/CPU0:router(config-ext)# 66:60001,
RP/0/RSP0/CPU0:router(config-ext)# 77:70001,
RP/0/RSP0/CPU0:router(config-ext)# 88:80001,
RP/0/RSP0/CPU0:router(config-ext)# 99:90001,

RP/0/RSP0/CPU0:router(config-ext)# 100.100.100.1:153
RP/0/RSP0/CPU0:router(config-ext)# end-set
```

extcommunity soo is-empty

To determine if a Border Gateway Protocol (BGP) route has any Site-of-Origin (SoO) extended communities associated with it, use the **extcommunity soo is-empty** command in route-policy configuration mode.

extcommunity soo is-empty

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo is-empty** command as a conditional expression within an **if** statement to check if a BGP SoO route has extended community attributes associated with it.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

The **is-empty** operator takes no arguments and evaluates to true if the route has no SoO extended community attributes associated with it.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, if a route has no SoO extended communities associated with it, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-is-empty-example
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo is-empty then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

extcommunity soo matches-any

To match any element of a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity soo matches-any** command in route-policy configuration mode.

extcommunity soo matches-any {*extcommunity-set-name* | *inline-extcommunity-set* | *parameter*}

Syntax Description

| | |
|--------------------------------|---|
| <i>extcommunity-set-name</i> | Name of a SoO extended community set. |
| <i>inline-extcommunity-set</i> | Inline SoO extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo matches-any** command as a conditional expression within an **if** statement to match elements of an extended community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-any** operator evaluates as true if at least one extended community in the route matches an extended community specification in the named or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, an SoO extended community set named extcomm-soo and a parameterized route policy named my-extcommunity-set-example(\$tag,\$ip) are defined.

The condition route policy named extcommunity soo matches-any is used in an if statement in this policy. If it evaluates to true, the local preference value is set to 100.

If it evaluates to false, the SoO extended community is evaluated using an inline set. If it evaluates to true, the local preference value is set to 200.

If it evaluates to false, the SoO extended community is evaluated using a different inline set. If it evaluates to true, the local preference value is set to 300.

If it evaluates to false, the SoO extended community is evaluated using a different inline set. If it evaluates to true, the local preference value is set to 400.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo extcomm-soo
RP/0/RSP0/CPU0:router(config-ext)# 66:60001,
RP/0/RSP0/CPU0:router(config-ext)# 77:70001,
RP/0/RSP0/CPU0:router(config-ext)# 88:80001,
RP/0/RSP0/CPU0:router(config-ext)# 99:90001,
RP/0/RSP0/CPU0:router(config-ext)# 100.100.100.1:153
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy my-extcommunity-set-example($tag,$ip)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo matches-any extcomm-soo then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-any (10:20, 10:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any ($ip:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif extcommunity soo matches-any (2.3.4.5:$tag)
then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 400
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|--|
| extcommunity rt matches-any, on page 1574 | Matches any element of a BGP RT extended community set. |
| extcommunity soo matches-every, on page 1590 | Matches every element of a BGP SoO extended community set. |

extcommunity soo matches-every

To match every element of a Border Gateway Protocol (BGP) Site-of-Origin (SoO) extended community set, use the **extcommunity soo matches-every** command in route-policy configuration mode.

extcommunity soo matches-every {*extcommunity-set-name*|*inline-extcommunity-set*|*parameter*}

Syntax Description

| | |
|--------------------------------|---|
| <i>extcommunity-set-name</i> | Name of a SoO extended community set. |
| <i>inline-extcommunity-set</i> | Inline SoO extended community set. The inline extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **extcommunity soo matches-every** command as a conditional expression within an **if** statement to match every element of a SoO extended community set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A simple condition using the **matches-every** operator evaluates as true if every extended community value in the extended community attribute for the route matches at least one element of the extended community set or inline set. If no extended community in the route matches any of the specifications in the named or inline set, then this simple condition evaluates to false. Likewise, when there is no extended community at all in the route, the condition evaluates to false.

Matching an extended community in the route to a specification in a named or an inline set is intuitive. In inline sets, extended community specifications may be parameterized, in which case the relevant matching is done when the value of the parameter has been supplied.

Examples

In the following example, an extended community set named my-extcomm-rt-set and a parameterized route policy named extcommunity-matches-every-example(\$as, \$tag) are defined. The condition extcommunity soo matches-every is used in an if statement in this policy and if it evaluates to true, the local-preference value is set to 100. If it evaluates to false, the extended community is evaluated using an inline set. If that condition evaluates to true, the local-preference value is set to 200. If it evaluates to false, the local-preference value is set to 300.

```
RP/0/RSP0/CPU0:router(config)# extcommunity-set soo my-extcomm-rt-set
RP/0/RSP0/CPU0:router(config-ext)# 10:20,
RP/0/RSP0/CPU0:router(config-ext)# 10:30,
RP/0/RSP0/CPU0:router(config-ext)# 10:40
RP/0/RSP0/CPU0:router(config-ext)# end-set

RP/0/RSP0/CPU0:router(config)# route-policy extcommunity-matches-every-example($as, $tag)
RP/0/RSP0/CPU0:router(config-rpl)# if extcommunity soo matches-every my-extcomm-rt-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif extcommunity soo matches-every (10:20, 10:$tag,
$as:30) then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 200
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 300
RP/0/RSP0/CPU0:router(config-rpl-elseif)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|--|
| extcommunity soo matches-any, on page 1588 | Matches any element of a BGP SoO extended community set. |

if

To decide which actions or dispositions should be taken for a given route, use the **if** command in route-policy configuration mode.

if *conditional-expression* **then** *action-statement* [*action-statement*] [**elseif** *conditional-expression* **then** *action-statement* [*action-statement*]] [**else** *action-statement* [*action-statement*]] **endif**

Syntax Description

| | |
|-------------------------------|---|
| <i>conditional-expression</i> | Expression to decide which actions or dispositions should be taken for the given route. |
| then | Executes an action statement if the if condition is true. |
| elseif | Strings together a sequence of tests. |
| else | Executes an action statement if the if condition is false. |
| endif | Ends the if statement. |
| <i>action-statement</i> | Sequence of operations that modify a route. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.2.0 | Support was added for Apply Condition Policies that allow the usage of a route-policy in an "if" statement of another route-policy. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **if** command uses a conditional expression to decide which actions or dispositions should be taken for a given route. [Table 172: Conditional Expressions, on page 1593](#) lists the conditional expressions.

An action statement is a sequence of operations that modify a route, most of which are distinguished by the **set** keyword. In a route policy, these operations can be grouped. [Table 173: Action Statements, on page 1595](#) lists the action statements.

Apply Condition policies allow usage of a route-policy in an "if" statement of another route-policy.

```
Route-policy policy_name
If apply policyA and apply policyB then
Set med 100
Else if not apply policyD then
Set med 200
Else
Set med 300
Endif
End-policy
```

Table 172: Conditional Expressions

| Command | Description |
|---|---|
| as-path in, on page 1526 | Matches the AS path of a route to an AS path set. The AS path is a sequence of autonomous system numbers traversed by a route. |
| as-path is-local, on page 1528 | Determines if the router (or another router within this autonomous system or confederation) originated the route. |
| as-path length, on page 1530 | Performs a conditional check based on the length of the AS path. |
| as-path neighbor-is, on page 1532 | Tests the autonomous system number or numbers at the head of the AS path against a sequence of one or more integral values or parameters. |
| as-path originates-from, on page 1535 | Tests an AS path against the AS sequence beginning with the AS number that originated a route. |
| as-path passes-through, on page 1538 | Tests to learn if the specified integer or parameter appears anywhere in the AS path or if the sequence of integers and parameters appears. |
| as-path unique-length, on page 1543 | Performs specific checks based on the length of the AS path. |
| community is-empty, on page 1545 | Learns if a route has community attributes associated with it. |
| community matches-any, on page 1546 | Matches any element of a community set. |
| community matches-every, on page 1548 | Matches every element of a community set. |
| destination in, on page 1558 | Matches a destination entry in a named prefix set or inline prefix set. |
| extcommunity rt is-empty, on page 1572 | Learns if a route has RT extended community attributes associated with it. |
| extcommunity rt matches-any, on page 1574 | Matches elements of an RT extended community set. |

| Command | Description |
|--|--|
| extcommunity rt matches-every, on page 1576 | Matches every element of an RT extended community set. |
| extcommunity rt matches-within, on page 1578 | Matches at least one element of a Border Gateway Protocol (BGP) route target (RT) extended community set. |
| extcommunity soo is-empty, on page 1586 | Learns if a route has SoO extended community attributes associated with it. |
| extcommunity soo matches-any, on page 1588 | Matches elements of an SoO extended community set. |
| extcommunity soo matches-every, on page 1590 | Matches every element of an SoO extended community set. |
| local-preference, on page 1599 | Specifies BGP local-preference attribute |
| med, on page 1601 | Compares the MED to an integer value or a parameterized value. |
| next-hop in, on page 1603 | Compares the next-hop associated with the route to data contained in either a named or an inline prefix set. |
| orf prefix in, on page 1605 | Matches a prefix in a prefix set or an inline prefix set. |
| origin is, on page 1607 | Tests the value of the origin attribute. |
| path-type is, on page 1611 | Tests the path type. |
| protocol, on page 1620 | Checks if a protocol is installing the route. |
| rd in, on page 1622 | Compares the RD associated with the route to data contained in either a named or an inline RD set. |
| rib-has-route, on page 1628 | Checks if a route is in the RIB. |
| route-has-label, on page 1630 | Checks if a route has a Multiprotocol Label Switching (MPLS) label. |
| route-type is, on page 1633 | Compares route types when redistribution is being performed into BGP, OSPF, or IS-IS. |
| source in, on page 1798 | Tests the source of the route against the data in either a named or an inline prefix set. |
| tag, on page 1802 | Matches a specific tag value. |

| Command | Description |
|---|---|
| vpn-distinguisher is , on page 1806 | Compares the VPN distinguisher against a specified value. |

Table 173: Action Statements

| Command | Description |
|---|---|
| abort (RPL) , on page 1520 | Discards a route policy definition and returns to global configuration mode. |
| add , on page 1522 | Adds an offset to an existing value. |
| apply , on page 1524 | Executes a parameterized or an unparameterized policy from within another policy. |
| delete community , on page 1554 | Deletes community values from a community list in a route. |
| delete extcommunity rt , on page 1556 | Deletes extended community values from an extended community list in a route. |
| done , on page 1560 | Accepts this route with no further processing |
| drop , on page 1562 | Drops a route. |
| end-policy , on page 1568 | Ends the definition of a route policy and exits route-policy configuration mode. |
| pass , on page 1609 | Signifies that even though the route has not been modified, the user wants to continue executing in the policy block. |
| prepend as-path , on page 1618 | Prepends the AS path with additional autonomous system numbers. |
| replace as-path , on page 1626 | Replaces a sequence of AS numbers or private AS numbers in the AS path with the configured local AS. |
| set community , on page 1641 | Sets the BGP community attribute. |
| set dampening , on page 1645 | Configures BGP route dampening. |
| set eigrp-metric , on page 1647 | Sets the Enhanced Interior Gateway Routing Protocol (EIGRP) metric value. |
| set extcommunity cost , on page 1649 | Replaces or adds the extended communities for a cost on the route. |

| Command | Description |
|--|---|
| set extcommunity rt , on page 1651 | Replaces or adds the extended communities for an RT on the route. |
| set ip-precedence , on page 1653 | Sets the IP precedence to classify packets. |
| set isis-metric , on page 1655 | Sets the IS-IS metric attribute value. |
| set label , on page 1656 | Sets the BGP label attribute value. |
| set level , on page 1659 | Configures the IS-IS level in which redistributed routes should be sent. |
| set local-preference , on page 1661 | Specifies a preference value for the autonomous system path. |
| set med , on page 1663 | Sets the MED value. |
| set metric-type (IS-IS) , on page 1665 | Controls whether IS-IS treats the metric as an internal or external metric. |
| set metric-type (OSPF) , on page 1667 | Controls whether OSPF treats the cost as a Type 1 or Type 2 metric. |
| set next-hop , on page 1669 | Replaces the next-hop associated with a given route. |
| set origin , on page 1671 | Changes the origin attribute. |
| set ospf-metric , on page 1673 | Sets an OSPF protocol metric attribute value. |
| set qos-group (RPL) , on page 1676 | Sets the QoS group to classify packets. |
| set rib-metric , on page 1678 | Sets a RIB metric attribute value for a table policy. |
| set rip-metric , on page 1679 | Sets RIP metric attributes. |
| set rip-tag , on page 1681 | Sets route tag attribute. |
| set tag , on page 1686 | Sets the tag attribute. |
| set traffic-index , on page 1688 | Sets the traffic index attribute. |
| set weight , on page 1692 | Sets the weight value for BGP routes. |
| suppress-route , on page 1800 | Indicates that a given component of an aggregate should be suppressed, that is, not advertised. |
| unsuppress-route , on page 1804 | Indicates that a given component of an aggregate should be unsuppressed. |

| Command | Description |
|--|-----------------------------------|
| set vpn-distinguisher , on page 1690 | Sets the VPN distinguisher value. |

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, any route whose AS path is in the set as-path-set-1 is dropped:

```
RP/0/RSP0/CPU0:router(config-rpl)# if as-path in as-path-set-1 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The contents of the **then** clause may be an arbitrary sequence of action statements.

The following example shows an **if** statement with two action statements:

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 42
RP/0/RSP0/CPU0:router(config-rpl-if)# prepend as-path 73 5
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The **if** command also permits an **else** clause to be executed if the expression is false, as follows:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 200 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set community (12:34) additive
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# set community (12:56) additive
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The routing policy language (RPL) also provides syntax using the **elseif** command to string together a sequence of tests, as shown in the following example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 10
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif med eq 200 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 60
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif med eq 250 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# set local-preference 110
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# set local-preference 0
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

The statements within an **if** statement may themselves be **if** statements, as shown in this example:

```
RP/0/RSP0/CPU0:router(config-rpl)# if community matches-any (12:34, 56:78) then
RP/0/RSP0/CPU0:router(config-rpl-if)# if med eq 150 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0/CPU0:router(config-rpl)#
```

The policy configuration shown sets the value of the local preference attribute to 100 on any route that has a community value of 12:34 or 56:78 associated with it. However, if any of these routes has a Multi Exit Discriminator (MED) value of 150, then each route with both the community value of 12:34 or 56:78 and a MED of 150 is dropped.

local-preference

To compare the local-preference attribute of a BGP route to an integer value or a parameterized value, use the local-preference command in route-policy configuration mode.

local-preference {**eq**|**is**|**ge**|**le**} {*number*|*parameter*}

Syntax Description

| | |
|---|---|
| eq is ge le | Equal to; exact match; greater than or equal to; less than or equal to. |
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **local-preference** command as a conditional expression within an **if** statement to compare the local-preference attribute to an integer value or a parameterized value.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The MED is a 32-bit unsigned integer. The **eq** operation compares the local-preference to either a static value or a parameterized value passed to a parameterized policy for equality with that value. A greater than or equal to comparison can also be done with the **ge** operator, and a less than or equal to comparison can be performed using the **le** operator.

Examples

The following example shows that if the **local-preference** is 10, local-preference is set to 100:

```
RP/0/RSP0RP0/CPU0:router(config-rpl)# if local-preference eq 10 then
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# set weight 100
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# endif
```

```
RP/0/RSP0RP0/CPU0:router(config-rpl)#
```

med

To compare the Multi Exit Discriminator (MED) to an integer value or a parameterized value or compare the MED attribute of a BGP route to an integer value, use the **med** command in route-policy configuration mode.

```
med {eq| is| ge| le} {number|parameter}
```

Syntax Description

| | |
|--------------------------|---|
| eq is ge le | Equal to; exact match; greater than or equal to; less than or equal to. |
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **med** command as a conditional expression within an **if** statement to compare the MED to an integer value or a parameterized value.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The MED is a 32-bit unsigned integer. The **eq** operation compares the MED to either a static value or a parameterized value passed to a parameterized policy for equality with that value. A greater than or equal to comparison can also be done with the **ge** operator, and a less than or equal to comparison can be performed using the **le** operator.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows that if the **med** commands match, the local preference is set to 100:

```
RP/0/RSP0/CPU0:router(config-rpl)# if med eq 10 then  
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 100  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

next-hop in

To compare the next-hop associated with the route to data contained in either an inline or a named prefix set, use the **next-hop in** command in route-policy configuration mode.

next-hop in {*prefix-set-name*| *inline-prefix-set*| *parameter*}

Syntax Description

| | |
|--------------------------|---|
| <i>prefix-set-name</i> | Name of a prefix set. |
| <i>inline-prefix-set</i> | Inline prefix set. The inline prefix set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **next-hop in** command as a conditional expression within an **if** statement to compare the next-hop associated with the route to data contained in either an inline or a named prefix set. The result is true if any value in the prefix set matches the next-hop of the route. A comparison that refers to a named prefix set that has no elements in it returns false.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The next-hop is an IPv4 address entered as a dotted-decimal or an IPv6 address entered as a colon-separated hexadecimal.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows that if the **next-hop in** commands match, the local preference is set to 100

```
RP/0/RSP0/CPU0:router(config-rpl)# if next-hop in some-prefix-set then
RP/0/RSP0/CPU0:router(config-rpl-if)# if next-hop in (10.0.0.5, fe80::230/64) then
RP/0/RSP0/CPU0:router(config-rpl-if)# set local-preference 0
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

orf prefix in

To configure an outbound route filter (ORF), use the **orf prefix in** command in route-policy configuration mode.

```
orf prefix in {prefix-set-name| inline-prefix-set}
```

Syntax Description

| | |
|--------------------------|---|
| <i>prefix-set-name</i> | Name of a prefix set. |
| <i>inline-prefix-set</i> | Inline prefix set. The inline prefix set must be enclosed in parentheses. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **orf prefix in** command to match a prefix in a prefix set or an inline prefix set.

This command takes either a named prefix set or an inline prefix set value as an argument. It returns true if the destination NLRI matches any entry in the prefix set. An attempt to match destination using a prefix set that is defined but contains no elements returns false.

This command is used in the context of the orf route-policy attach point in BGP. The destination of a route is also known in Border Gateway Protocol (BGP) as its network-layer reachability information (NLRI). It comprises a prefix value and a mask length. The routing policy language (RPL) provides one operation on prefixes, testing them for matching against a list of prefix-match specifications using the **in** operator.

Examples

In the following example, the prefix set orfpreset1 and the route policy named orfpolicy are defined. Next, the orfpolicy is applied to the neighbor orf attach point.

If the prefix of the route matches any of the prefixes specified in orfpreset1 (211.105.1.0/24, 211.105.5.0/24, 211.105.11.0/24), then the prefix is dropped. If the prefix matches in(211.105.3.0/24, 211.105.7.0/24, 211.105.13.0/24), then the prefix is accepted. In addition to this inbound filtering, BGP sends these prefix

entries to the upstream neighbor indicating a permit or deny so that the neighbor can make the same filter updates.

```
RP/0/RSP0/CPU0:router(config)# prefix-set orfpreset1
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.1.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.5.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 211.105.11.0/24
RP/0/RSP0/CPU0:router(config-pfx)# end-set
!
!
RP/0/RSP0/CPU0:router(config)# route-policy orfpolicy
RP/0/RSP0/CPU0:router(config-rpl)# if orf prefix in orfpreset1 then
RP/0/RSP0/CPU0:router(config-rpl-if)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# if orf prefix in (211.105.3.0/24, 211.105.7.0/24,
211.105.13.0/24) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
!
!
RP/0/RSP0/CPU0:router(config)# router bgp 2
RP/0/RSP0/CPU0:router(config-bgp)# neighbor 1.1.1.1
RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 3
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-nbr-af)# orf route-policy orfpolicy
```

Related Commands

| Command | Description |
|------------|---|
| orf | Specifies BGP ORF and inbound filtering criteria. |

origin is

To match a specific origin type, use the **origin is** command in route-policy configuration mode.

origin is {**igp**|**egp**|**incomplete**|*parameter*}

Syntax Description

| | |
|-------------------|---|
| igp | Specifies Interior Gateway Protocol. |
| egp | Specifies Exterior Gateway Protocol. |
| incomplete | Specifies that Border Gateway Protocol (BGP) first learned the route by means other than BGP or Interior Gateway Protocol (IGP); for example, the route is learned through configuration. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **origin is** command as a conditional expression within an **if** statement to test the value of the origin attribute.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The origin of a BGP route is an enumeration; it is **igp** , **egp** , or **incomplete** .

This command can be parameterized.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the origin is tested within an **if** statement to learn if it is either **igp** or **egp** :

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp or origin is egp then
```

In the following example, a parameter is used to match a specific origin type:

```
RP/0/RSP0/CPU0:router(config)# route-policy bar($origin)
RP/0/RSP0/CPU0:router(config-rpl)# if origin is $origin then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 20
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

pass

To pass a route for further processing, use the **pass** command in route-policy configuration mode.

pass

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **pass** command to signify that even though this route has not been modified, the user wants to continue executing in this policy block.



Note

The **pass** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

When a policy block has finished executing, any route that has been modified in this policy block or has received a pass disposition in this policy block passes the policy and execution finishes for that policy. If this policy block is applied from within another policy block and the route is either passed or modified, then execution continues in the policy block that applied this policy block.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to accept the route unconditionally without modifying it:

```
RP/0/RSP0/CPU0:router(config-rpl)# pass
```

This example accepts the route unconditionally, without modifying it, if the destination is in prefix-set permitted:

```
RP/0/RSP0/CPU0:router(config-rpl)# if destination in permitted then  
RP/0/RSP0/CPU0:router(config-rpl-if)# pass  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

path-type is

To match path types, use the **path-type is** command in route-policy configuration mode.

path-type is {**ibgp**|**ebgp**| *parameter*}

| Syntax Description | | |
|--------------------|------------------|--|
| | ibgp | Specifies an internal BGP path. |
| | ebgp | Specifies an external BGP path. |
| | <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **path-type is** command as a conditional expression within an **if** statement to match path types.



Note For a list of all conditional expressions available within an **if** statement, see the **if** command.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, if the path is an external BGP path the route is accepted:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_A
RP/0/RSP0/CPU0:router(config-rpl)# if path-type is ebgp then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
```

path-type is

```
RP/0/RSP0/CPU0:router(config-rpl-if)# else  
RP/0/RSP0/CPU0:router(config-rpl-else)# drop  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

policy-global

To define global parameters and enter global parameter configuration mode, use the **policy-global** command in global configuration mode. To remove global parameters, use the **no** form of this command.

policy-global

no policy-global

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **policy-global** command to define global parameters and enter global parameter configuration mode. RPL supports the definition of systemwide global parameters that can be used inside a policy definition. The global parameter values can be used directly inside a policy definition similar to the local parameters of parameterized policy. When a parameterized policy has a parameter name “collision” with a global parameter name, parameters local to policy definition take precedence, effectively ‘masking off’ global parameters. In addition, a validation mechanism is in place to prevent the deletion of a particular global parameter if it is referred by any policy. For more information on global parameters and parameterization, see the *Implementing Routing Policy on Cisco ASR 9000 Series Router Cisco IOS XR Software* module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples The following example shows how to configure global parameters:

```
RP/0/RSP0/CPU0:router(config)# policy-global
RP/0/RSP0/CPU0:router(config-rp-gl)# glbpathtype 'ebgp'
```

```
RP/0/RSP0/CPU0:router(config-rp-gl)# glbtag '100'
RP/0/RSP0/CPU0:router(config-rp-gl)# end-global
```

In the following example, the *globalparam* argument makes use of the global parameters *gblpathtype* and *glbtag* defined above and is defined for a nonparameterized policy:

```
RP/0/RSP0/CPU0:router(config)# route-policy globalparam
RP/0/RSP0/CPU0:router(config-rpl)# if path-type is $gblpathtype then
RP/0/RSP0/CPU0:router(config-rpl)# set tag $glbtag
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|---|
| end-global, on page 1567 | Ends the definition of global parameters. |

prefix-set

To enter prefix set configuration mode and define a prefix set for contiguous and non-contiguous set of bits, use the **prefix-set** command in global configuration mode. To remove a named prefix set, use the **no** form of this command.

prefix-set *name*

no prefix-set *name*

Syntax Description

| | |
|-------------|-----------------------|
| <i>name</i> | Name of a prefix set. |
|-------------|-----------------------|

Command Default

None

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 5.1.1 | The command was modified to accept IPv4 and IPv6 address and address mask to define a prefix set. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **prefix-set** command to enter prefix set configuration mode and define a prefix set.

A prefix set is a comma-separated list of prefix match specifications. It holds IPv4 or IPv6 prefix match specifications, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. The address is a standard four-part, dotted-decimal numeric IPv4 address or a colon-separated hexadecimal IPv6 address. The mask length, if present, is a nonnegative decimal integer in the range from 0 to 32 for IPv4 prefixes or 0 to 128 for IPv6 prefixes following the address and separated from it by a slash. The optional minimum matching length follows the address and optional mask length and is expressed as the keyword **ge** (mnemonic for **greater than or equal to**), followed by a nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. The optional maximum matching length follows the rest and is expressed by the keyword **le** (mnemonic for **less than or equal to**), followed by yet another nonnegative decimal integer in the range from 0 to 32 for IPv4 or 0 to 128 for IPv6. A syntactic shortcut for specifying an exact length for prefixes to match is the **eq** keyword, mnemonic for **equal to**.

If a prefix match specification has no mask length, then the default mask length is 32 for IPv4 or 128 for IPv6. The default minimum matching length is the mask length. If a minimum matching length is specified, then

the default maximum matching length must be less than 32 for IPv4 prefixes or 128 for IPv6 prefixes. Otherwise, if neither a minimum nor maximum length is specified, the default maximum length is the mask length.

A prefix set is a list of prefix match specifications. It holds IPv4 or IPv6 prefix match specifications, each of which has two parts: an address and a mask. The address and mask is a standard dotted-decimal IPv4 or colon-separated hexadecimal IPv6 address. The prefix set allows the specifying of contiguous and non-contiguous set of bits that must be matched in any route. The set of bits to be matched are provided in the form of a mask in which a binary 0 means a mandatory match and a binary 1 means a 'do not match' condition.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows a prefix set named legal-ipv4-prefix-examples:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv4-prefix-examples
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.1.1,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.2.0/24,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.3.0/24 ge 28,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.4.0/24 le 28,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.5.0/24 ge 26 le 30,
RP/0/RSP0/CPU0:router(config-pfx)# 10.0.6.0/24 eq 28
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

The first element of the prefix set matches only one possible value, 10.0.1.1/32 or the host address 10.0.1.1. The second element matches only one possible value, 10.0.2.0/24. The third element matches a range of prefix values, from 10.0.3.0/28 to 10.0.3.255/32. The fourth element matches a range of values, from 10.0.4.0/24 to 10.0.4.240/28. The fifth element matches prefixes in the range from 10.0.5.0/26 to 10.0.5.252/30. The sixth element matches any prefix of length 28 in the range from 10.0.6.0/28 through 10.0.6.240/28.

The following prefix set consists entirely of invalid prefix match specifications:

```
RP/0/RSP0/CPU0:router(config)# prefix-set INVALID-PREFIX-EXAMPLES
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.1.1 ge 16,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.2.1 le 16,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.3.0/24 le 23,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.4.0/24 ge 33,
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.5.0/25 ge 29 le 28
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

Neither the minimum length nor the maximum length is legal without a mask length. The maximum length must be at least the mask length. The minimum length must be less than 32, the maximum length of an IPv4 prefix. The maximum length must be equal to or greater than the minimum length.

The following example shows a valid IPv6 prefix set named legal-ipv6-prefix-examples:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv6-prefix-examples
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:1::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:2::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:3::/64,
RP/0/RSP0/CPU0:router(config-pfx)# 2001:0:0:4::/64
RP/0/RSP0/CPU0:router(config-pfx)# end-set
```

Examples

This example shows a prefix set named legal-ipv4-prefix:

```
RP/0/RSP0/CPU0:router(config)# prefix-set legal-ipv4-prefix  
RP/0/RSP0/CPU0:router(config-pfx)# 10.1.1.1 0.255.0.255  
RP/0/RSP0/CPU0:router(config-pfx)# 10.2.2.2 0.0.0.0  
RP/0/RSP0/CPU0:router(config-pfx)# 10.3.3.3 255.255.255.254  
RP/0/RSP0/CPU0:router(config-pfx)# 10.4.4.4 255.255.255.255
```

In the above example, the command defines the prefix-set named legal-ipv4-prefix. The first element specifies to match all routes having 10 in first octet and 1 in third octet. The second element matches all routes having prefix as 10.2.2.2 (that is, matches all conditions). The third element matches all routes having odd numbers in the last octets and the fourth element matches all routes with any prefix.

prepend as-path

To prepend the AS path with additional autonomous system numbers, use the **prepend as-path** command in route-policy configuration mode.

prepend as-path {*as-number*|*parameter*|**most-recent**} [*number*|*parameter*]

Syntax Description

| | |
|--------------------|--|
| <i>as-number</i> | Autonomous system number to prepend to the path. <ul style="list-style-type: none"> • Range for 2-byte Autonomous system numbers (ASNs) is 1 to 65535. • Range for 4-byte Autonomous system numbers (ASNs) in asplain format is 1 to 4294967295. • Range for 4-byte Autonomous system numbers (ASNs) in asdot format is 1.0 to 65535.65535. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| most-recent | Specifies that the most recent autonomous system number should be prepended. |
| <i>number</i> | (Optional) Number of times the autonomous system number should be prepended. Range is 1 to 63. |

Command Default

The default *number* is 1.

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 3.9.0 | Asplain format for 4-byte Autonomous system numbers notation was supported. |

Usage Guidelines

Use the **prepend as-path** command to prepend the AS path with additional autonomous system numbers.



Note

The **prepend as-path** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command can take one or two arguments. The first argument (either a number or parameter) is the autonomous system number to prepend to the path. The optional second argument (either a number or parameter) is the number of times the autonomous system number should be prepended.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to prepend the autonomous system number 666.1 to the AS path three times:

```
RP/0/RSP0/CPU0:router(config-rpl)# prepend as-path 666.1 3
```

The following example shows how to prepend the autonomous system number 666.0 to the AS path one time:

```
RP/0/RSP0/CPU0:router(config-rpl)# prepend as-path 666.0 1
```

protocol

To check the protocol that installs the route, use the **protocol** command in route-policy configuration mode.

protocol {**in** | (*protocol-set*) | **is** | *protocol-name*}

Syntax Description

| | |
|-----------------------------------|--|
| in (<i>protocol-set</i>) | <p>Specifies a member of a set. The <i>protocol-set</i> argument accepts the following keywords within parentheses:</p> <ul style="list-style-type: none"> • bgp —Border Gateway Protocol (BGP) • connected —Connected routes • eigrp —Enhanced Interior Gateway Routing Protocol (EIGRP) • isis —ISO Intermediate System-to-Intermediate System (IS-IS) • ospf —Open Shortest Path First (OSPF) • ospfv3 —Open Shortest Path First version 3 (OSPFv3) • rip —Routing Information Protocol (RIP) • static —Static routes <p>Keywords must be separated by a comma.</p> |
| is <i>protocol-name</i> | <p>Specifies a single protocol name, and accepted keywords are similar to the <i>protocol-set</i> argument.</p> |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **protocol** command as a conditional expression within an if statement to specify a protocol to install a route.

Use the **in** keyword to determine if a protocol listed in the *protocol-set* is the originator of the route being filtered.

Use the **is** keyword to determine if *protocol-name* is an exact match.

**Note**

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to use the **protocol** command as a conditional expression within if statements:

```
RP/0/RSP0/CPU0:router(config)# route-policy rip1
RP/0/RSP0/CPU0:router(config-rpl)# if protocol in (connected, static) then
RP/0/RSP0/CPU0:router(config-rpl-if)# add rip-metric 2
RP/0/RSP0/CPU0:router(config-rpl-if)# elseif protocol is bgp 1 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# add rip-metric 3
RP/0/RSP0/CPU0:router(config-rpl-elseif)# elseif protocol is ospf 2 then
RP/0/RSP0/CPU0:router(config-rpl-elseif)# add rip-metric 4
RP/0/RSP0/CPU0:router(config-rpl-elseif)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# add rip-metric 5
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy

RP/0/RSP0/CPU0:router(config)# router rip
RP/0/RSP0/CPU0:router(config-rip)# interface GigabitEthernet0/1/0/1
RP/0/RSP0/CPU0:router(config-rip-if)# route-policy rip1 out
```

rd in

To compare the route distinguisher (RD) associated with the route to RDs contained in either a named or an inline RD set, use the **rd in** command in route-policy configuration mode.

```
rd in {rd-set-name| inline-rd-set| parameter}
```

Syntax Description

| | |
|----------------------|---|
| <i>rd-set-name</i> | Name of an RD set. |
| <i>inline-rd-set</i> | Inline RD set. The inline RD set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rd in** command as a conditional expression within an **if** statement to match a destination entry in a named prefix set or inline prefix set.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command takes either a named RD set or an inline RD set value as an argument. The condition returns true if the destination entry matches any entry in the RD set or inline RD set. An attempt to match an RD using an RD set that is defined but contains no elements returns false.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows the **rd in** command with an inline RD set value as an argument:

```
RP/0/RSP0/CPU0:router(config)# route-policy
RP/0/RSP0/CPU0:router(config-rpl)# if rd in (128.1.0.0/16:100) then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

rd-set

To define a route distinguisher (RD) set and enter RD configuration mode, use the **rd-set** command in global configuration mode.

rd-set *name*

no rd-set *name*

Syntax Description

| | |
|------|------------------------------|
| name | Name of an RD community set. |
|------|------------------------------|

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rd-set** command to create a set with RD elements and enter RD configuration mode. An RD set is a 64-bit value prepended to an IPv4 address to create a globally unique Border Gateway Protocol (BGP) VPN IPv4 address.



Note

For *m*, the mask length is supported.

You can define RD values with the following commands:

- *a.b.c.d/m:**—BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2/24.0.*.
- *a.b.c.d/m:n*—BGP VPN RD in IPv4 format with a mask. For example, 10.0.0.2/24:666.
- *a.b.c.d.** —BGP VPN RD in IPv4 format with a wildcard character. For example, 10.0.0.2.*.
- *a.b.c.d:n*—BGP VPN RD in IPv4 format. For example, 10.0.0.2:666.
- *asn:**—BGP VPN RD in ASN format with a wildcard character. For example, 10002.*.
- *asn:n*—BGP VPN RD in ASN format. For example, 10002:666.
- *x.y:**—BGP VPN RD in 4-byte ASN format with a wildcard character. For example, 10002.101.*.

- *x.y:n*—BGP VPN RD in 4-byte ASN format. For example, 10002.101:666.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to create an RD set called `my_rd_set`:

```
RP/0/RSP0/CPU0:router(config)# rd-set my_rd_set
RP/0/RSP0/CPU0:router(config-rd)# 172.16.0.0/16:*
RP/0/RSP0/CPU0:router(config-rd)# 172.17.0.0/16:100
RP/0/RSP0/CPU0:router(config-rd)# 192:*
RP/0/RSP0/CPU0:router(config-rd)# 192:100
RP/0/RSP0/CPU0:router(config-rd)# end-set
```

replace as-path

To replace a sequence of AS numbers or private AS numbers in the AS path with the configured local AS number, use the **replace as-path** command in route-policy configuration mode.

```
replace as-path {[as-number-list parameter]} private-as
```

Syntax Description

| | |
|-----------------------|---|
| <i>as-number-list</i> | (Optional) Sequence of AS numbers to replace. The sequence must be enclosed in single quotes (' '). You can use 2-byte or 4-byte AS numbers. <ul style="list-style-type: none"> The 2-byte value is entered as a 16-bit unsigned decimal value. The range is 0 to 65535. The 4-byte value is entered as two 16-bit unsigned decimal values separated by a period. The range is 1.0 to 65535.65535. |
| <i>parameter</i> | (Optional) Parameter name. The parameter name must be preceded with a "\$." |
| private-as | Matches within the BGP private AS range. Range is from 64512 to 65534. |

Command Default

None.

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 3.7.2 | This command was introduced. |
| Release 4.1.0 | This command was supported on ASR 9000 Ethernet Line Card (Cisco ASR 9000's A9K-SIP-700). |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **replace as-path** command to replace a sequence of AS numbers or private AS numbers in the AS path with the local AS numbers. For example, if the AS path is '67 65534 100 65533 5 78 89 90' and the local AS number is 900, then:

```
replace as-path '5 78'
```

replaces '5 78' in the AS path with 900 (from the local AS), and the new path would be '67 65534 100 65533 900 89 90'.

Consider following statement:

```
replace as-path private-as
```

Because 65534 and 65533 are within the private AS range, they are replaced with 900. The path is '67 900 100 900 5 78 89 90'. The length of the path remains the same.

The **replace as-path** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.



Caution

The **replace as-path** command changes the AS path content which can lead to routing loops.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to use the **replace as-path** command to replace AS numbers in the AS path:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '90 78 45 $asnum'
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path private-as
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '9.9 7.89 14.15 $asnum'
RP/0/RSP0/CPU0:router(config-rpl)# replace as-path '9 89 14.15 $asnum'
```

rib-has-route

To check if a route listed in the prefix set exists in the Routing Information Base (RIB), use the **rib-has-route** command in route-policy configuration mode.

rib-has-route in {*prefix-set-name*| *inline-prefix-set*| *parameter*}

Syntax Description

| | |
|--------------------------|---|
| <i>prefix-set-name</i> | Name of a prefix set. |
| <i>inline-prefix-set</i> | Inline prefix set. The inline prefix set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If routes are active, then they are advertised. Routes are considered active if they are already installed in the Routing Information Base (RIB).

The prefix sets used in the **rib-has-route** command contain two match specifications. The first is where an exact route match is requested (for example, 10.10.0.0/16 will match exactly one route) and the second is where a route match or any more-specific route match is allowed (for example, 10.10.0.0/16 le 32 will match the 10.10.0.0/16 route and any longer prefix).

Use the **rib-has-route** command as a conditional expression within an **if** statement to check if there is an active route with a specific prefix contained in the RIB. If the statement reveals an active route that meets that criteria, additional actions are executed.

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, an **if** statement is used to learn if a route contained in a prefix set 10.10.0.0/16 is in the RIB:

```
RP/0/RSP0/CPU0:router(config-rpl)# if rib-has-route in (10.10.0.0/16 ge 16) then  
RP/0/RSP0/CPU0:router(config-rpl-if)# pass  
RP/0/RSP0/CPU0:router(config-rpl-if)# endif  
RP/0/RSP0/CPU0:router(config-rpl)#
```

route-has-label

To check if there is a Multiprotocol Label Switching (MPLS) label in a route during redistribution, use the **route-has-label** command in route-policy configuration mode.

route-has-label

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-has-label** command as a conditional expression within an **if** statement to check if there is an MPLS label in a route during redistribution.

For a list of all conditional expressions available within an **if** statement, see the **if** command.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, an **if** statement learns if an MPLS label is present in a route:

```
RP/0/RSP0/CPU0:router(config-rpl)# if route-has-label then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

route-policy (RPL)

To define a route policy and enter route-policy configuration mode, use the **route-policy** command in global configuration mode. To remove a policy definition, use the **no** form of this command.

route-policy *name* [(*parameter1*, *parameter2*, . . . , *parameterN*)]

no route-policy *name* (*parameter1*, *parameter2*, . . . , *parameterN*)

Syntax Description

| | |
|------------------|---|
| <i>name</i> | Name of a route policy. |
| <i>parameter</i> | (Optional) Parameter name. The parameter name must be preceded with a "\$." The <i>parameters</i> must be enclosed in parenthesis "()". |

Command Default

No default behavior or values

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-policy** command to define a route policy and enter route-policy configuration mode.

Policy definitions create named bundles of policy statements. A policy definition consists of the **route-policy** command followed by a name, a group of policy statements, and the **end-policy** command.

The policy name serves as a handle for binding the policy to protocols.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows a simple policy named drop-everything that drops any route it encounters:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-everything
RP/0/RSP0/CPU0:router(config-rpl)# drop
```

```
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Policies may also refer to other policies such that common blocks of policy can be reused. This reference to other policies is accomplished by using the **apply** command. The following is a simple example:

```
RP/0/RSP0/CPU0:router(config)# route-policy drop-as-1234
RP/0/RSP0/CPU0:router(config-rpl)# if as-path passes-through '1234' then
RP/0/RSP0/CPU0:router(config-rpl-if)# apply check-communities
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# pass
RP/0/RSP0/CPU0:router(config-rpl-else)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

The **apply** command indicates that the policy check-communities should be executed if the route under consideration passed through autonomous system 1234 before it was received. If so, the communities of the route are checked, and based on the findings the route may be accepted unmodified, accepted with changes, or dropped.

Related Commands

| Command | Description |
|---|--|
| end-policy , on page 1568 | Ends the definition of a route policy. |

route-type is

To match route types when redistribution is being performed into Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), or Integrated Intermediate System-to-Intermediate System (IS-IS), use the **route-type is** command in route-policy configuration mode.

```
route-type is {local| interarea| internal| type-1| type-2| level-1| level-2| parameter}
```

Syntax Description

| | |
|------------------|---|
| local | Uses a local value to match locally generated BGP routes. |
| interarea | Uses an interarea value to match IS-IS interarea routes. |
| internal | Uses an internal value to match OSPF intra- and interarea routes. |
| type-1 | Uses a Type 1 value to match Type 1 OSPF routes. |
| type-2 | Uses a Type 2 value to match Type 2 OSPF routes. |
| level-1 | Uses a Level 1 value to match Level 1 IS-IS routes. |
| level-2 | Uses a Level 2 value to match Level 2 IS-IS routes. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **route-type is** command as a conditional expression within an **if** statement to compare route types when redistribution is being performed into BGP, OSPF, or IS-IS.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The valid keywords are **local**, **internal**, **interarea**, **type-1**, **type-2**, **level-1**, and **level-2**. A parameterized value that fills in one of these values may also be used. The **local** value is used to match locally generated BGP routes. The **internal** value is used to match OSPF intra- and interarea routes. The **type-1** and **type-2** values are used to match Type 1 and Type 2 OSPF external routes. The **level-1**, **level-2**, and **interarea** values are used to match IS-IS routes of those respective types.

Because the route type is a matching operator, it appears in conditional clauses of **if** and **then** statements.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, non-local routes are dropped:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_A
RP/0/RSP0/CPU0:router(config-rpl)# if route-type is local then
RP/0/RSP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0/CPU0:router(config-rpl-else)# drop
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

rpl editor

To set the default routing policy language (RPL) editor, use the **rpl editor** command in global configuration mode.

rpl editor {**nano**|**emacs**|**vim**}

Syntax Description

| | |
|--------------|--|
| nano | Sets the default RPL editor to GNU nano. |
| emacs | Sets the default RPL editor to EMACS. |
| vim | Sets the default RPL editor to VIM. |

Command Default

The Nano editor is the default.

Command Modes

Global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the default RPL editor is set to Nano:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl editor nano
```

In the following example, the default RPL editor is set to EMACS:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl editor emacs
```

In the following example, the default RPL editor is set to VIM:

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# rpl editor vim
```

rpl maximum

To configure system limits on the routing policy subsystem, use the **rpl maximum** command in global configuration mode.

rpl maximum {*lines*| *policies*} *number*

Syntax Description

| | |
|-------------------------------|---|
| lines <i>number</i> | Configures the number of lines of configuration limit. Range is from 1 to 131072. |
| policies <i>number</i> | Configures the number of policies limit. Range is from 1 to 5000. |

Command Default

lines *number* : 65536
policies *numbers* : 3500

Command Modes

global configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **rpl maximum** command to configure system limits on the routing policy subsystem. As such, **rpl maximum** configuration lines do not appear as statements within a routing policy. This command places resource limits on the routing policy subsystem. Use the **rpl maximum** command to configure the maximum number of lines of configuration and number of policies.

The number of lines of configuration includes the beginning and ending statements, for example, **route-policy** and **end-policy**. Each line of configuration for sets is also counted.

A line of configuration is counted only once; it is not counted each time it is used. Similarly, any multiple use of policy in an apply statement counts only as one policy.

A user can change the default values for lines and policies but cannot exceed the maximum value, nor can the value for lines and policies be configured lower than the number of lines or policies that are currently configured.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the maximum number of RPL system limits are modified:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# rpl maximum lines 50
RP/0/RSP0/CPU0:router(config)# rpl maximum policies 6
```

Related Commands

| Command | Description |
|--|--|
| show rpl maximum, on page 1749 | Displays the maximum limits for lines of configuration and number of policies. |

set aigp-metric

To set originating prefixes with accumulated interior gateway protocol (AiGP) attribute information, use the **set aigp-metric** command in route-policy configuration mode.

```
set aigp-metric {igp-cost | value}
```

| Syntax Description | Parameter | Description |
|--------------------|-----------------|--|
| | igp-cost | Specifies the internal routing protocol cost. |
| | <i>value</i> | Specifies the aigp-metric value. 32-bit decimal number. Range is 0-4294967295. |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.0.0 | This command was introduced. |
| | Release 5.0.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples The following example shows how to set the aigp metric as the igp cost for route-policy aigp_policy:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy aigp_policy
RP/0/RSP0/CPU0:router(config-rpl)# set aigp-metric igp-cost
```

Related Commands

| Command | Description |
|---------------------------------|---|
| aigp | Enables sending and receiving of accumulated interior gateway protocol (AiGP) attribute per neighbor. |
| aigp send-cost-community | Sends accumulated interior gateway protocol (AiGP) value in cost community. |

set community

To set the Border Gateway Protocol (BGP) community attributes in a route, use the **set community** command in route-policy configuration mode.

```
set community {community-set-name| inline-community-set| parameter} [additive]
```

Syntax Description

| | |
|-----------------------------|---|
| <i>community-set-name</i> | Community set name. |
| <i>inline-community-set</i> | Inline community set. The inline community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| additive | (Optional) Adds communities to communities in the route. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set community** command to set the BGP community attribute.



Note

The **set community** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Communities are 32-bit values carried in BGP routes. Each route may have zero or more communities in an unordered list.

Use this command to replace the communities in a route or add to them using the optional **additive** keyword.

As with the other community forms that support inline sets, either or both 16-bit portions of the community can be parameterized. Likewise, the names of the well-known communities **internet** (0:0), **no-advertise** (65535:65281), **no-export** (65535:65282), and **local-AS** (65535:65283) can also be used. In an inline community set, each 16-bit portion can also be specified as the **peeras** to express the AS number of the

neighbor from which the route was received. If the neighbor AS employs a 4-byte ASN, the IANA-assigned 16-bit value 23456 (AS_TRANS) is used as **peeras** instead.

Without the **additive** keyword, any existing communities (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all communities already present in the route be maintained and the list of communities be added to them.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following are incomplete configuration examples using the **set community** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24)
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, $as:24, $as:$tag)
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, internet) additive
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, $as:24) additive
RP/0/RSP0/CPU0:router(config-rpl)# set community (10:24, peeras:24) additive
```

set core-tree

To set a Multicast Distribution Tree (MDT) type, use the **set core-tree** command in route-policy configuration mode.

```
set core-tree {gre-rosen| mldp-inband| mldp-partitioned-mp2mp| mldp-partitioned-p2mp| mldp-rosen|
rsvp-te-partitioned-p2mp| parameter}
```

Syntax Description

| | |
|---------------------------------|--|
| gre-rosen | Specifies the IP GRE Rosen core MDT type |
| mldp-inband | Specifies the MLDP InBand core MDT type |
| mldp-partitioned-mp2mp | Specifies theMLDP Partitioned MP2MP core MDT type |
| mldp-partitioned-p2mp | Specifies the MLDP Partitioned P2MP core MDT type |
| mldp-rosen | Specifies the MLDP Rosen core MDT type |
| rsvp-te-partitioned-p2mp | Specifies the RSVP TE core core MDT type |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

None

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.1.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|--------------|-------------|
| route-policy | read, write |

Examples

In this example, the Multicast Distribution Tree type is set to IP GRE Rosen core:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#route-policy policy_mdt_type
RP/0/RSP0/CPU0:router(config-rpl)#set core-tree gre-rosen
```

set dampening

To configure Border Gateway Protocol (BGP) route dampening, use the **set dampening** command in route-policy configuration mode.

```
set dampening {halflife {minutes| parameter}| max-suppress {minutes| parameter}| reuse {seconds| parameter}| suppress {penalty-units| parameter}| others default}
```

Syntax Description

| | |
|--------------------------------------|--|
| halflife <i>minutes</i> | Specifies the time (in minutes) after which a penalty is decreased. After the route has been assigned a penalty, the penalty is decreased by half after the half-life period. The process of reducing the penalty happens every 5 seconds. Range is 1 to 45 minutes. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| max-suppress <i>minutes</i> | Specifies the maximum time (in minutes) a route can be suppressed. Range is 1 to 20000. If the half-life value is allowed to default, the maximum suppress time defaults to 60 minutes. |
| reuse <i>seconds</i> | Unsuppresses a route if the penalty for flapping the route decreases enough to fall below the configured value (in seconds). The process of unsuppressing routes occurs at 10-second increments. Range is 1 to 20000. |
| suppress <i>penalty-units</i> | Specifies a penalty of 1000 each time a route flaps. When a route penalty exceeds the configured limit, it is suppressed. Range is 1 to 20000. |
| others default | If all four keyword values are not specified in the command, then the command <i>must</i> end with others default . This designation indicates that any keyword not defined is set to its default. |

Command Default

half-life : 15 minutes
max-suppress : 60 minutes (four times the half-life)
reuse : 750 seconds
suppress : 2000 penalty units

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The BGP protocol supports route dampening using an exponential backoff algorithm. The algorithm is controlled by setting the four supported BGP values: half-life, max-suppress, reuse, and suppress. Use the **set dampening** command to configure BGP route dampening.

**Note**

The **set dampening** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A value for at least one of the four keywords must be set. If the **set dampening** command defines values for three or fewer of the supported keywords, then the configuration must end with the **others default**, which indicates that any keyword value not defined in the command is set to its default value.

The keywords may appear in the command in any order.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following examples, the half-life is set to 20 minutes and the maximum suppress time is set to 90 minutes. Each command must end with **others default** because three or fewer keywords are defined.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife 20 others default
RP/0/RSP0/CPU0:router(config-rpl)# set dampening max-suppress 90 others default
```

In this example, all four keywords are defined, which means the command does not use **others default**.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife 15 max-suppress 60 reuse 750
suppress 2000
```

The following command is invalid because it is missing **others default**.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening reuse 700
```

In the following example, the parameters are used.

```
RP/0/RSP0/CPU0:router(config-rpl)# set dampening halflife $p1 suppress $p4 reuse $p3
max-suppress $p2
```

set eigrp-metric

To set Enhanced Interior Gateway Routing Protocol (EIGRP) route metrics, use the **set eigrp-metric** command in route-policy configuration mode.

set eigrp-metric *bandwidth delay reliability loading mtu*

Syntax Description

| | |
|--------------------|--|
| <i>bandwidth</i> | Minimum bandwidth of the route in kilobits per second. Range is 1 to 4294967295. |
| <i>delay</i> | Route delay in tens of microseconds. Delay is 1 or any positive number that is a multiple of 39.1 nanoseconds. Range is 1 to 4294967295. |
| <i>reliability</i> | Likelihood of successful packet transmission expressed as a number between 0 and 255. The value 255 means 100 percent reliability; 0 means no reliability. |
| <i>loading</i> | Effective bandwidth of the route expressed as a number from 1 to 255 (255 is 100 percent loading). |
| <i>mtu</i> | Minimum maximum transmission unit (MTU) size of the route in bytes. Range is from 1 to 65535. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

You can use the **add** command to further offset an existing EIGRP metric value.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the EIGRP metrics are adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set eigrp-metric 1400 120 250 100 1500
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|--------------------------------------|
| add , on page 1522 | Adds an offset to an existing value. |

set extcommunity cost

To set the Border Gateway Protocol (BGP) cost extended community attributes, use the **set extcommunity cost** command in route-policy configuration mode.

```
set extcommunity cost {cost-extcommunity-set-name| cost-inline-extcommunity-set| parameter} [additive]
```

Syntax Description

| | |
|-------------------------------------|---|
| <i>cost-extcommunity-set-name</i> | Cost extended community set name. |
| <i>cost-inline-extcommunity-set</i> | Inline cost extended community set. The inline cost extended community set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| additive | (Optional) Adds extended communities for cost to extended communities in the route. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set extcommunity cost** command to either replace the extended communities on the route or add to them using the optional **additive** keyword. Cost community is an extended community used to tie break the best path selection process in BGP so as to have a localized custom decision for packet forwarding. The extended community format defines generic points of insertion (POI) that influence the decision at different points of the bestpath algorithm.



Note

The **set extcommunity cost** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

As with the other extended community forms that support inline sets, either or both portions of the community can be parameterized. Similarly to regular communities, the **additive** keyword can be used to signify adding these extended communities to those that are already present, as opposed to replacing them. Without the

additive keyword, any existing extended communities for cost (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all extended communities for cost already present in the route be maintained and the set of extended communities be added to them. Well-known communities include internet, local-AS, no-advertise, and no-export.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following are incomplete configuration examples using the **set extcommunity cost** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (IGP:10:20)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (Pre-Bestpath:33:44)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity cost (IGP:11:21)
```

set extcommunity rt

To set the Border Gateway Protocol (BGP) route target (RT) extended community attributes, use the **set extcommunity rt** command in route-policy configuration mode.

```
set extcommunity rt {rt-extended-community-set-name| rt-inline-extended-community-set} parameter} additive
```

| Syntax Description | | |
|--------------------|---|---|
| | <i>rt-extended-community-set-name</i> | Route target extended community set name. |
| | <i>rt-inline-extended-community-set</i> | Inline route target extended community set. The inline route target extended community set must be enclosed in parentheses. |
| | <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| | additive | (Optional) Adds extended communities for an RT to extended communities in the route. |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set extcommunity rt** command to either replace the extended communities on the route or add to them using the optional **additive** keyword.



Note

The **set extcommunity rt** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

As with the other extended community forms that support inline sets, either or both portions of the community can be parameterized. Similarly to regular communities, the **additive** keyword can be used to signify adding these extended communities to those that are already present, as opposed to replacing them.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples

The following are incomplete configuration examples using the **set extcommunity rt** command:

```
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, $as:24, $as:$tag)
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, internet) additive
RP/0/RSP0/CPU0:router(config-rpl)# set extcommunity rt (10:24, $as:24) additive
```

Without the **additive** keyword, any existing extended communities for cost (other than the well-known communities) are removed and replaced with the given communities. The **additive** keyword specifies that all extended communities for cost already present in the route be maintained and the list of extended communities be added to them.

set ip-precedence

To set the IP precedence, use the **set ip-precedence** command in route-policy configuration mode.

```
set ip-precedence {number|parameter}
```

Syntax Description

number Value of the precedence. The precedence value can be a number from 0 to 7:

- 7 —network (set packets with network control precedence)
- 6 —internet (set packets with internetwork control precedence)
- 5 —critical (set packets with critical precedence)
- 4 —flash-override (set packets with flash override precedence)
- 3 —flash (set packets with flash precedence)
- 2 —immediate (set packets with immediate precedence)
- 1 —priority (set packets with priority precedence)
- 0 —routine (set packets with routine precedence)

parameter Parameter name. The parameter name must be preceded with a "\$."

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set ip-precedence** command to set the IP precedence to classify packets. This command is supported at the BGP table-policy attachpoint. Prefixes are marked for subsequent processing in the forwarding plane. After QoS Policy Propagation through Border Gateway Protocol (BGP) (QPPB) is enabled on an interface, corresponding traffic shaping and policing is completed using packet classification based on the IP precedence or QoS group ID. See *Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide* for information on QPPB.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| route-policy | read, write |

Examples

This example shows how use **set ip-precedence** command:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1  
RP/0/RSP0/CPU0:router(config-rpl)# set ip-precedence 3  
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set isis-metric

To set the Intermediate System-to-Intermediate System (IS-IS) metric attribute value, use the **set is-is metric** command in route-policy configuration mode.

```
set isis-metric {number|parameter}
```

| Syntax Description | |
|--------------------|--|
| <i>number</i> | 24-bit integer number. Range is from 0 to 16777215. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set isis-metric** command to set the IS-IS metric attribute value for routes that are redistributed into IS-IS.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, the IS-IS metric attribute value is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set isis-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set label

To set the Border Gateway Protocol (BGP) label attribute value, use the **set label** command in route-policy configuration mode.

set label {**explicit-null**|**implicit-null**| *parameter*}

Syntax Description

| | |
|----------------------|--|
| explicit-null | Sets the label to the well-known explicit value of 0. |
| implicit-null | Sets the label to the well-known implicit value of 3. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set label** command in a route policy at the allocate label attachpoint to set the label to explicit-null or implicit-null based on deployment preference. During inter-AS operation, the ASBR sends some of its own loopbacks to other its peers and labels them either implicit null or explicit null.

Examples

The following example shows how to set the labels:

```
RP/0/RSP0/CPU0:router(config)# route-policy labelpolicy
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (206.141.1.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RSP0/CPU0:router(config-rpl)# elseif destination in (206.141.3.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# drop
RP/0/RSP0/CPU0:router(config-rpl)# elseif destination in (206.141.4.0/24) then
RP/0/RSP0/CPU0:router(config-rpl)# set label explicit-null
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set label-mode

To set the type of Border Gateway Protocol (BGP) label mode, use the **set label-mode** command in route-policy configuration mode. This command does not have a **no** form.

```
set label-mode {per-ce | per-vrf | per-prefix}
```

Syntax Description

| | |
|-------------------|--|
| per-ce | Specifies that the same label is used for all routes advertised from a unique customer edge (CE) peer or router. |
| per-vrf | Specifies that the same label is used for all routes advertised from a unique VRF. |
| per-prefix | Specifies that the same label is used for all routes advertised from a unique prefix. |

Command Default

Per-prefix label mode.

If a policy attached at label-mode attachpoint evaluates to pass and a **label mode** is not explicitly set, **per-prefix** is used as a default label mode.

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set label-mode** command in a route policy at the label-mode attachpoint to set the type of label mode to per-ce or per-vrf or per-prefix, based on deployment preference.

Task ID

| Task ID | Operation |
|--------------|-------------|
| route-policy | read, write |

Examples

This example shows how to set the type of label-mode to per-ce:

```
RP/0/RSP0/CPU0:router(config)# route-policy set_label_mode
```

```
RP/0/RSP0/CPU0:router(config-rpl)# set label-mode per-ce
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

This example shows how to set the type of label-mode to per-vrf:

```
RP/0/RSP0/CPU0:router(config)# route-policy set_label_mode
RP/0/RSP0/CPU0:router(config-rpl)# set label-mode per-vrf
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

This example shows how to set the type of label-mode to per-prefix:

```
RP/0/RSP0/CPU0:router(config)# route-policy set_label_mode
RP/0/RSP0/CPU0:router(config-rpl)# set label-mode per-prefix
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|--|
| route-policy (RPL), on page 1631 | Defines a route policy and enters route-policy configuration mode. |

set level

To configure the Intermediate System-to-Intermediate System (IS-IS) link-state packet (LSP) level advertised to redistributed routes, use the **set level** command in route-policy configuration mode.

```
set level {level-1| level-2| level-1-2| parameter}
```

Syntax Description

| | |
|------------------|---|
| level-1 | Specifies that redistributed routes are advertised in the Level 1 LSP of the router. |
| level-2 | Specifies that redistributed routes are advertised in the Level 2 LSP of the router. |
| level-1-2 | Specifies that redistributed routes are advertised in Level 1 and Level 2 LSPs of the router. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the IS-IS **set level** command to configure the LSP level advertised to redistributed routes.



Note

The **set level** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command supports parameterization of the **level** keyword.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the level is set to Level 2:

```
RP/0/RSP0/CPU0:router(config)# route-policy bgp_isis_redist
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (172.2.0.0/16 ge 16) then
RP/0/RSP0/CPU0:router(config-rpl)# set level level-2
RP/0/RSP0/CPU0:router(config-rpl)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set local-preference

To set the Border Gateway Protocol (BGP) local preference attribute in a route, use the **set local-preference** command in route-policy configuration mode.

set local-preference {*number*|*parameter*}

Syntax Description

| | |
|------------------|--|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

Default value is 100.

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set local-preference** command to specify a preference value for the autonomous system path. Local preference is a nontransitive (does not cross autonomous system boundaries) attribute and is the second metric considered in the BGP best path calculation (the highest local preference is chosen). Weight is the first metric evaluated for best path, but it is local to the router and propagates only to iBGP peers. See the *Implementing BGP on Cisco ASR 9000 Series Router Cisco IOS XR Software* module of the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide* for information on the BGP best path calculation.



Note

The **set local-preference** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The local preference is a 32-bit unsigned integer.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the local preference value is set to 10:

```
RP/0/RSP0/CPU0:router(config-rpl)# set local-preference 10
```

set med

To set the Border Gateway Protocol (BGP) Multi Exit Discriminator (MED) attribute, use the **set med** command in route-policy configuration mode.

```
set med {number|parameter} igp-cost {+| {number|parameter}} |-| {number|parameter}}| max-reachable}
```

Syntax Description

| | |
|----------------------|--|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| igp-cost | Sets the MED value to the cost for the Interior Gateway Protocol (IGP) route to resolve the next-hop of the BGP route. |
| + - | Sets the MED to the MED plus or minus a static offset. An integer or parameter must follow the plus or minus. |
| max-reachable | Sets the MED value to the maximum possible value of 4294967295. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set med** command to set the MED value, which is a 32-bit unsigned integer.



Note

The **set med** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command can take the following as argument values: an integer, a parameter, the **igp-cost** keyword, or a mathematical operator (either plus or minus) followed by an integer or a parameter. Setting the MED to the IGP cost is supported on outbound BGP policies only. The MED cannot be set to the IGP cost in policies applied to other BGP attach points.

The **max-reachable** keyword sets the MED to the maximum value while leaving the route reachable.

The plus or minus variants allow the user to set the MED to the MED plus or minus a static offset. The variants that allow a user to add or subtract offsets to the MED value are also range checked for underflow or overflow. If the value underflows as a result of subtraction, then the MED value is set to zero. If the value overflows, the value is set to 4294967295, which is the maximum value for MED. when MED is set to 4294967295, the route is unreachable.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following two examples show how to set the MED to a value that is either specified directly (using the integer 156) or passed to the policy as a parameter:

```
RP/0/RSP0/CPU0:router(config-rpl)# set med 156
RP/0/RSP0/CPU0:router(config-rpl)# set med $med_param
```

The following example shows how to instruct BGP to automatically set the MED value to the cost of the IGP route that resolves the next-hop of the BGP route:

```
RP/0/RSP0/CPU0:router(config-rpl)# set med igp-cost
```

set metric-type (IS-IS)

To configure the integrated Intermediate System-to-Intermediate System (IS-IS) metric type, use the **set metric-type** command in route-policy configuration mode.

```
set metric-type {internal| external| parameter}
```

| Syntax Description | | |
|--------------------|------------------|--|
| | internal | Sets metric type to internal. |
| | external | Sets the metric type to external. |
| | <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the IS-IS **set metric-type** command to control whether IS-IS treats the metric as an internal or external metric.



Note The **set metric-type** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

This command does not support parameterization.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples

In the following example, the IS-IS metric type is set to internal:

```
RP/0/RSP0/CPU0:router(config-rpl)# set metric-type internal
```

set metric-type (OSPF)

To control how Open Shortest Path First (OSPF) computes the cost for a route, use the **set metric-type** command in route-policy configuration mode.

```
set metric-type {type-1| type-2| parameter}
```

| Syntax Description | | |
|--------------------|------------------|---|
| | type-1 | Uses the cost set on the route plus the topology-related costs in the calculation for Type 1 metrics. |
| | type-2 | Uses only the cost set on the route in the calculation for Type 2 metrics. |
| | <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the OSPF **set metric-type** command to control whether OSPF treats the cost as a Type 1 or Type 2 metric.



Note The **set metric-type** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The value of Type 1 or Type 2 controls how OSPF computes the cost for this route. For Type 2 metrics, only the cost set on the route is used. For Type 1 metrics, the cost set on the route plus the topology-related costs are used in the calculation.

This command does not support parameterization.

Task ID

| Task ID | Operations |
|----------------|-------------------|
| route-policy | read, write |

Examples

In the following example, the OSPF metric type is set to Type 1:

```
RP/0/RSP0/CPU0:router(config-rpl)# set metric-type type-1
```

set next-hop

To replace the next-hop associated with a given route, use the **set next-hop** command in route-policy configuration mode.

```
set next-hop {ipv4-address| ipv6-address| peer-address| parameter| self} [destination-vrf]
```

Syntax Description

| | |
|------------------------|---|
| <i>ipv4-address</i> | Valid IPv4 address. |
| <i>ipv6-address</i> | Valid IPv6 address. |
| peer-address | Sets the next-hop to the IP address of the remote Border Gateway Protocol (BGP) peer. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| self | Sets itself as the next-hop. |
| destination-vrf | (Optional) Specifies that the next-hop of the route should be resolved in destination VRF context. This keyword is available when an IPv4 or IPv6 address or parameter is used. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set next-hop** command to replace the next-hop associated with a specific address.



Note

The **set next-hop** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Use the **set next-hop peer-address** command to set the next-hop to the address of the BGP neighbor, where this policy is attached.

The next-hop is a valid IPv4 address entered as a dotted decimal or an IPv6 address entered as a colon-separated hexadecimal.

It is not possible to use this command to set the BGP IPv6 link-local next-hop.

The **destination-vrf** keyword is used mainly in Layer 3 VPN networks when importing routes.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the next-hop is set to a valid IPv4 address:

```
RP/0/RSP0/CPU0:router(config-rpl) # set next-hop 10.0.0.5
```

In this example, the next-hop is set to a parameter value \$nexthop:

```
RP/0/RSP0/CPU0:router(config-rpl) # set next-hop $nexthop
```

In this example, the next-hop is set to a valid IPv4 address with a destination VRF context:

```
RP/0/RSP0/CPU0:router(config-rpl) # set next-hop 10.0.0.5 destination-vrf
```

set origin

To change the Border Gateway Protocol (BGP) origin attribute, use the **set origin** command in route-policy configuration mode.

```
set origin {igp|incomplete|egp|parameter}
```

Syntax Description

| | |
|-------------------|--|
| igp | Sets the origin type to Interior Gateway Protocol (IGP). |
| incomplete | Sets the origin type to incomplete. |
| egp | Sets the origin type to Exterior Gateway Protocol (EGP). |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set origin** command to change the origin attribute.



Note

The **set origin** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

The origin of a Border Gateway Protocol (BGP) route is **igp**, **egp**, or **incomplete**.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the origin attribute is set to EGP:

```
RP/0/RSP0/CPU0:router(config-rpl)# set origin egp
```

set ospf-metric

To set an Open Shortest Path First (OSPF) protocol metric attribute value, use the **set ospf-metric** command in route-policy configuration mode.

```
set ospf-metric {number|parameter}
```

| Syntax Description | | |
|--------------------|--|--|
| <i>number</i> | | Value assigned to a 24-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | | Parameter name. The parameter name must be preceded with a "\$." |

Command Default No default behavior or values

Command Modes Route-policy configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set ospf-metric** command to set the metric for routes that are redistributed into OSPF. The OSPF metric operator accepts either an integer value or a parameter.

| Task ID | Task ID | Operations |
|---------|--------------|-------------|
| | route-policy | read, write |

Examples In the following example, the OSPF metric attribute value is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set ospf-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set path-selection

To set path selection criteria and install or advertise the path for the Border Gateway Protocol, use the **set path-selection** command in route-policy configuration mode.

set path-selection {*backup number*| **group-best**| **all**| **best-path**} [**install**] [**multipath-protect**] [**advertise**]

Syntax Description

| | |
|--------------------------|---|
| backup | Specifies the BGP backup path. |
| <i>number</i> | Specifies the BGP backup path number. 3 bit decimal number. Range is 0-7. |
| group-best | Specifies the BGP group best path. |
| all | Specifies all BGP paths. |
| best-path | Specifies the BGP best path. |
| install | Installs the path. |
| multipath-protect | Installs and advertises the multipath protect. |
| advertise | Advertises the path. |

Command Default

None

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|---|
| Release 4.0.0 | This command was introduced. |
| Release 4.0.1 | The multipath-protect keyword was added. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|--------------|-------------|
| route-policy | read, write |

Examples

The following example shows how to set the path selection as **advertise backup path 3** for route-policy *path_selection_plcy*:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# route-policy path_selection_plcy
RP/0/RSP0/CPU0:router(config-rpl)# set path-selection backup 3 advertise
```

Related Commands

| Command | Description |
|-----------------------------------|--|
| additional-paths selection | Configures additional paths selection capability for a prefix |
| additional-paths send | Configures send capability of multiple paths for a prefix to the capable peers, |
| additional-paths receive | Configures receive capability of multiple paths for a prefix to the capable peers. |
| advertise best-external | Advertises the best-external path to the iBGP and route-reflector peers, |

set qos-group (RPL)

To set the quality of service (QoS) group, use the **set qos-group** command in route-policy configuration mode:

```
set qos-group {number|parameter}
```

Syntax Description

| | |
|------------------|--|
| <i>number</i> | QoS group ID. Range is from 0 to 31. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set qos-group** command to set the QoS group to classify packets.

This command is supported at the BGP table-policy attachpoint. Prefixes are marked for subsequent processing in the forwarding plane. After QoS Policy Propagation through Border Gateway Protocol (BGP) (QPPB) is enabled on an interface, corresponding traffic shaping and policing is completed using packet classification based on the IP precedence or QoS group ID. See the *Cisco ASR 9000 Series Aggregation Services Router Modular Quality of Service Configuration Guide* for information on QPPB.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

This example shows how to use **set qos-group** command:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set qos-group 12
```

```
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rib-metric

To set the Routing Information Base (RIB) metric attribute value for a table policy, use the **set rib-metric** command in route-policy configuration mode:

```
set rib-metric {number|parameter}
```

Syntax Description

| | |
|------------------|--|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rib-metric** command set the RIB metric attribute value for BGP routes.

Every route in the RIB has a metric associated with it, signifying the cost to reach a specific destination based on link characteristics. The **set rib-metric** command modifies the RIB metric while installing BGP routes into RIB, enabling the upgrading or downgrading of the BGP route installed in RIB.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the RIB metric attribute is set to 1000:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rib-metric 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rip-metric

To set Routing Information Protocol (RIP) metric attributes, use the **set rip-metric** command in route-policy configuration mode.

```
set rip-metric {number|parameter}
```

Syntax Description

| | |
|------------------|--|
| <i>number</i> | Value assigned to a 4-bit unsigned integer. Range is from 0 to 16. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rip-metric** command to set the cost attribute for routes that are redistributed into RIP.

You can use the **add** command to increment the RIP metric value.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the RIP metric number is adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rip-metric 10
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Related Commands

| Command | Description |
|--|--------------------------------------|
| add , on page 1522 | Adds an offset to an existing value. |

set rip-tag

To set a route tag attribute for Routing Information Protocol (RIP) routes, use the **set rip-tag** command in route-policy configuration mode.

```
set rip-tag {number|parameter}
```

Syntax Description

| | |
|------------------|--|
| <i>number</i> | Value assigned to a 16-bit unsigned integer. Range is from 0 to 65535. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set rip-tag** command to set the RIP tag attribute for routes that are redistributed into RIP. The RIP tag operator accepts either an integer value or a parameter.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the RIP tag is adjusted for route policy policy_1:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy_1
RP/0/RSP0/CPU0:router(config-rpl)# set rip-tag 1000
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

set rpf-topology

To set reverse-path forwarding (RPF) to any default or nondefault tables for particular sources and groups, use the **set rpf-topology** command in routing policy configuration mode.

```
set rpf-topology [vrf vrf-name] {ipv4| ipv6} {unicast| multicast| parameter} topology table-name
```

Syntax Description

| | |
|-------------------|--|
| ipv4 | [Optional] Specifies IPv4 address prefixes. |
| ipv6 | [Optional] Specifies IPv6 address prefixes. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| topology | Specifies the default or non-default topology table for the source or group. |
| <i>table-name</i> | Alphanumeric name string. |

Command Default

Default or current topology setting.

Command Modes

Routing policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|-----------|------------|
| multicast | read |

Examples

The following example shows how to execute the **set rpf-topology** command:

```
RP/0/RSP0/CPU0:router# config
RP/0/RSP0/CPU0:router(config)# route-policy green
RP/0/RSP0/CPU0:router(config-rpl)# set rpf-topology ipv6 multicast topology t12
```

The following example shows the use of **set rpf-topology** command in the context of creating an RPF for a topology table in multiple topologies:

```
route-policy mt4-p1
  if destination in (225.0.0.1, 225.0.0.11) then
    set rpf-topology ipv4 multicast topology t201
  elseif destination in (225.0.0.2, 225.0.0.12) then
    set rpf-topology ipv4 multicast topology t202
  elseif destination in (225.0.0.3, 225.0.0.13) then
    pass
  endif
end-policy
!

route-policy mt4-p3
  if destination in (225.0.0.8) then
    set rpf-topology ipv4 multicast topology t208
  elseif destination in (225.0.0.9) then
    set rpf-topology ipv4 multicast topology t209
  elseif destination in (225.0.0.10) then
    set rpf-topology ipv4 multicast topology t210
  else
    drop
  endif
end-policy
!
```

Related Commands

| Command | Description |
|--------------|---|
| rpf topology | Assigns a route policy in PIM to select a reverse-path forwarding (RPF) topology. |

set spf-priority

To set OSPF Shortest Path First (SPF) priority, use the set spf-priority command in route-policy configuration mode.

```
set spf-priority {critical| high| medium}
```

Syntax Description

| | |
|-----------------|--------------------------------|
| critical | Sets critical priority for SPF |
| high | Sets high priority for SPF |
| medium | Sets medium priority for SPF |

Command Default

None

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|--------------|-------------|
| route-policy | read, write |

Examples

This example sets SPF priority as critical:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#route-policy policy_spf_priority
RP/0/RSP0/CPU0:router(config-rpl)#set spf-priority critical
```

Related Commands

| Command | Description |
|----------------------------|---|
| spf prefix-priority (OSPF) | Prioritizes OSPFv2 prefix installation into the global Routing Information Base (RIB) during Shortest Path First (SPF) run. |

set tag

To set the tag attribute, use the **set tag** command in route-policy configuration mode.

```
set tag {number|parameter}
```

Syntax Description

| | |
|------------------|---|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is from 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set tag** command to set the tag attribute.



Note

The **set tag** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Tags are routing-protocol independent 32-bit integers that can be associated with a given route in the Routing Information Base (RIB).

For the Border Gateway Protocol (BGP), the tag attribute can be set only at the table-policy attach point.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the tag attribute is set to 10:

```
RP/0/RSP0/CPU0:router(config-rpl)# set tag 10
```

In this example, the tag attribute is set to a parameter value \$tag_param:

```
RP/0/RSP0/CPU0:router(config-rpl)# set tag $tag_param
```

set traffic-index

To set the traffic index attribute, use the **set traffic-index** command in route-policy configuration mode.

```
set traffic-index {number|parameter|ignore}
```

Syntax Description

| | |
|------------------|---|
| <i>number</i> | Integer value assigned to the traffic index attribute. Range is 1 to 63. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |
| ignore | Specifies that Border Gateway Protocol (BGP) policy accounting is not done. |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set traffic-index** command to set the traffic index attribute.



Note

The **set traffic-index** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Traffic index is a special attribute for BGP. It is used as an index to a set of counters that are maintained by forwarding hardware. It is also used to track packet and byte counters that are forwarded using routes with specific attributes. These counters can be enabled and disabled on an individual interface basis.

The traffic index attribute can be set only at the table-policy attach point, and can take a value from 1 to 63, or a value of **ignore**. If the traffic index is set to **ignore**, then BGP policy accounting is not done. Parameterization of this value is also supported.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, a policy is created in which the traffic index is set to 10 for all routes that originated in autonomous system 1234:

```
RP/0/RSP0RP0/CPU0:router(config)# route-policy count-as-1234
RP/0/RSP0RP0/CPU0:router(config-rpl)# if as-path originates-from '1234' then
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# set traffic-index 10
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# else
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# pass
RP/0/RSP0RP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0RP0/CPU0:router(config-rpl)# end-policy
```

This policy could then be attached using the BGP **table-policy** command. The counters could then be enabled on various interfaces with the appropriate commands.

set vpn-distinguisher

To change the Border Gateway Protocol (BGP) VPN distinguisher attribute, use the **set vpn-distinguisher** command in route-policy configuration mode.

```
set vpn-distinguisher {number|parameter}
```

Syntax Description

| | |
|------------------|---|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set vpn-distinguisher** command to change the VPN distinguisher attribute.



Note

The **set origin** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A VPN distinguisher is used in Layer 3 VPN networks for enhanced individual VPN control and to avoid route target mapping at AS boundaries in inter-AS VPN networks. Route target extended communities are removed at neighbor outbound, and the VPN distinguisher value is applied on the BGP route as an extended community. When the route is received on a neighboring router in another AS, the VPN distinguisher is removed and mapped to a route target extended community.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the VPN distinguisher attribute is set to 456:

```
RP/0/RSP0/CPU0:router(config-rpl)# set vpn-distinguisher 456
```

set weight

To set the weight value for Border Gateway Protocol (BGP) routes, use the **set weight** command in route-policy configuration mode.

set weight {*number*|*parameter*}

Syntax Description

| | |
|------------------|---|
| <i>number</i> | Number assigned to the weight value for BGP routes. Weight is 16 bits. Range is 0 to 65535. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **set weight** command to set the weight value for BGP routes.



Note

The **set weight** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

A weight is a value that can be applied to a route to override the BGP local preference. This is not a BGP attribute announced to BGP peer routers. RPL can be used to set the weight value.

Given two BGP routes with the same network layer reachability information (NLRI), a route with a higher weight is selected, no matter what the values of other BGP attributes may be. However, weight only has significance on the local router. It is not sent from one BGP speaker to another, even within the same autonomous system.

On Cisco routers, if a BGP route is sourced by the local router, its weight is automatically set to 32768; if the BGP route is learned from another router, its weight is automatically set to 0. Thus, by default, locally sourced routes are preferred over BGP learned routes.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the weight of the route is set to 10 and then to a parameter value \$weight_param:

```
RP/0/RSP0/CPU0:router(config-rpl)# set weight 10
RP/0/RSP0/CPU0:router(config-rpl)# set weight $weight_param
```

show rpl

To display system-wide RPL configuration, use the **show rpl** command in EXEC mode.

show [**running-config**] **rpl** [**maximum** {**lines** *configuration-limit*| **policies** *policies-limit*}| **editor** {**emacs**| **nano**| **vim**}]

Syntax Description

| | |
|---|---|
| running-config | (Optional) Displays configuration-limit argument. |
| maximum | (Optional) Displays the maximum number of lines of configuration and number of policies. |
| lines <i>configuration-limit</i> | (Optional) Displays the number of lines to which configuration is limited. Range is 1 to 131072. The <i>configuration-limit</i> argument is available if the running-config keyword is specified. |
| policies <i>policies-limit</i> | (Optional) Displays the limit on the number of policies. Range is 1 to 5000. The <i>configuration-limit</i> argument is available if the running-config keyword is specified. |
| editor | (Optional) Specifies the default RPL editor. This keyword is available if the running-config keyword is specified. |
| emacs | (Optional) Displays the default RPL editor to Micro Emacs. |
| nano | (Optional) Displays the default RPL editor to nano. |
| vim | (Optional) Displays the default RPL editor to Vim. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

The following shows the output of the **show running-config rpl** command:

```
RP/0/RSP0/CPU0:router# show running-config rpl

extcommunity-set rt ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any com_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
```

Related Commands

| Command | Description |
|--|--|
| show rpl maximum, on page 1749 | Displays the maximum limits for lines of configuration and number of policies. |

show rpl active as-path-set

To display the AS path sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active as-path-set** command in EXEC mode.

show rpl active as-path-set [detail]

| | |
|---------------------------|---|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for active AS path sets. |
|---------------------------|---|

Command Default No default behavior or values

Command Modes EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active as-path-set** command to display all AS path sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt_ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!

```

Given this sample configuration, the **show rpl active as-path-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl active as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are ACTIVE
-----
as_path_set_ex1

```

Related Commands

| Command | Description |
|--|--|
| show rpl active community-set, on page 1699 | Displays the community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active extcommunity-set, on page 1702 | Displays the extended community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the route policies that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the prefix sets that are referenced by at least one policy that is being used at an attach point. |

show rpl active community-set

To display the community sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active community-set** command in EXEC mode.

show rpl active community-set [detail]

| | |
|---------------------------|---|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for active community sets. |
|---------------------------|---|

| | |
|------------------------|-------------------------------|
| Command Default | No default behavior or values |
|------------------------|-------------------------------|

| | |
|----------------------|------|
| Command Modes | EXEC |
|----------------------|------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active community-set** command to display all community sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
 !
 !
 neighbor 10.0.101.3
```

show rpl active community-set

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt_ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!

route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!

```

Given this sample configuration, the **show rpl active community-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl active community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are ACTIVE

```

 comm_set_ex1

Related Commands

| Command | Description |
|--|---|
| show rpl active as-path-set, on page 1696 | Displays the AS path sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active extcommunity-set, on page 1702 | Displays the extended community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the route policies that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the prefix sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active rd-set, on page 1708 | Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point. |

show rpl active extcommunity-set

To display the extended community sets for cost, route target (RT), and Site-of-Origin (SoO) that are referenced by at least one route policy used at an attach point, use the **show rpl active extcommunity-set** command in EXEC mode.

show rpl active extcommunity-set [**cost**] [**rt**] [**soo**] [**detail**]

Syntax Description

| | |
|---------------|--|
| cost | (Optional) Displays all extended community cost sets. |
| rt | (Optional) Displays all extended community RT sets. |
| soo | (Optional) Displays all extended community SoO sets. |
| detail | (Optional) Displays the content of the object and all referenced objects for active extended community sets. |

Command Default

All extended community sets are displayed.

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active extcommunity-set** command to display all extended community sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
```

```

address-family ipv4 unicast
!
neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
  route-policy policy_1 in
!
!
neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
  route-policy policy_2 in
!
!
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt_ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!

route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!

```

Given this sample configuration, the **show rpl active extcommunity-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached

UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are ACTIVE
-----
ext_comm_set_rt_ex1
```

Related Commands

| Command | Description |
|---|---|
| show rpl active as-path-set, on page 1696 | Displays the AS path sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active community-set, on page 1699 | Displays the community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the route policies that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the prefix sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active rd-set, on page 1708 | Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point. |

show rpl active prefix-set

To display the prefix sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active prefix-set** command in EXEC mode.

show rpl active prefix-set [detail]

| | |
|---------------------------|--|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for active prefix sets. |
|---------------------------|--|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active prefix-set** command to display all prefix sets that are in use in the system and referenced either directly or indirectly at a policy attach point.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
 !
 !
 neighbor 10.0.101.3
```

show rpl active prefix-set

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!

community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
extcommunity-set rt_ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!

route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!

RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail

prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_655--$',
  ios-regex '^_65501_$'
end-set
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!

```

The following example displays active prefix sets:

```

RP/0/RSP0/CPU0:router# show rpl active prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are ACTIVE
-----

```

```
prefix_set_1
```

Related Commands

| Command | Description |
|--|---|
| show rpl active as-path-set, on page 1696 | Displays the AS path sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active community-set, on page 1699 | Displays the community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active extcommunity-set, on page 1702 | Displays the extended community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl route-policy attachpoints, on page 1771 | Displays the route policies that are referenced by at least one policy that is being used at an attach point. |
| show rpl active rd-set, on page 1708 | Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point. |

show rpl active rd-set

To display the route distinguisher (RD) sets that are referenced by at least one policy that is being used at an attach point, use the **show rpl active rd-set** command in EXEC mode.

show rpl active rd-set [detail]

Syntax Description

| | |
|---------------|---|
| detail | (Optional) Displays the content of the object and all referenced objects for active route policies. |
|---------------|---|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active rd-set** command to display all RD sets that are in use in the system and that are referenced either directly or indirectly at a policy attach point.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
rd-set rdset1
  10:151,
  100.100.100.1:153,
  100.100.100.62/31:63
end-set
!
rd-set rdset2
  10:152,
  100.100.100.1:154,
  100.100.100.62/31:89
```

```

end-set
!
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
end-policy
!
router bgp 10
  bgp router-id 10.0.0.1
  address-family vpnv4 unicast
neighbor 10.10.10.1
  remote-as 10
  address-family ipv4 unicast
  route-policy rdsetmatch in
!
!

```

Given this sample configuration, the **show rpl active rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl active rd-set
```

```
ACTIVE -- Referenced by at least one policy which is attached INACTIVE -- Only referenced
by policies which are not attached UNUSED -- Not attached (directly or indirectly) and not
referenced
```

```
The following rd-sets are ACTIVE
```

```
-----
rdset1
rdset2
```

Related Commands

| Command | Description |
|--|--|
| show rpl active as-path-set, on page 1696 | Displays the AS path sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active community-set, on page 1699 | Displays the community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active extcommunity-set, on page 1702 | Displays the extended community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the prefix sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the route policies that are referenced by at least one policy that is being used at an attach point. |

show rpl active route-policy

To display the route policies that are referenced by at least one policy that is being used at an attach point, use the **show rpl active route-policy** command in EXEC mode.

show rpl active route-policy [detail]

| | |
|---------------------------|---|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for active route policies. |
|---------------------------|---|

Command Default No default behavior or values

Command Modes EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl active route-policy** command to display all policies that are in use in the system and that are referenced either directly or indirectly at a policy attach point.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!

```

Given this sample configuration, the **show rpl active route-policy** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl active route-policy

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following policies are (ACTIVE)
-----
policy_1
policy_2

```

Related Commands

| Command | Description |
|--|--|
| show rpl active as-path-set, on page 1696 | Displays the AS path sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active community-set, on page 1699 | Displays the community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active extcommunity-set, on page 1702 | Displays the extended community sets that are referenced by at least one policy that is being used at an attach point. |
| show rpl active prefix-set, on page 1705 | Displays the prefix sets that are referenced by at least one policy that is being used at an attach point. |

| Command | Description |
|--|---|
| show rpl active rd-set, on page 1708 | Displays the route distinguisher sets that are referenced by at least one policy that is being used at an attach point. |

show rpl as-path-set

To display the contents of AS path sets, use the **show rpl as-path-set** command in EXEC mode.

show rpl as-path-set [*name*] **states** [**brief**]

| Syntax Description | |
|--------------------|--|
| name | (Optional) Name of the AS path set. |
| states | (Optional) Displays all unused, inactive, and active states. |
| brief | (Optional) Limits the display to a list of the names of all AS path sets without their configurations. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of all AS path sets without their configurations.

| Task ID | Task ID | Operations |
|---------|--------------|------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
```

```

if (as-path in as_path_set_ex1) then
  set community (10:333) additive
endif
end-policy

```

Given this sample configuration, the **show rpl as-path-set as_path_set_ex1** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1

as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set

```

Related Commands

| Command | Description |
|---|---|
| show rpl community-set, on page 1721 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1729 | Displays the configuration of a named extended community set. |
| show rpl route-policy, on page 1768 | Displays the configuration of a named route policy. |
| show rpl prefix-set, on page 1753 | Displays the configuration of a named prefix set. |

show rpl as-path-set attachpoints

To display all of the policies used at an attach point that reference the named AS path set, use the **show rpl as-path-set attachpoints** command in EXEC mode.

show rpl as-path-set *name* **attachpoints**

| | | |
|---------------------------|-------------|-------------------------|
| Syntax Description | <i>name</i> | Name of an AS path set. |
|---------------------------|-------------|-------------------------|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl as-path-set attachpoints** command to display all policies used at an attach point that reference the named set either directly or indirectly.

The AS path set name is required.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

show rpl as-path-set attachpoints

```

remote-as 12
address-family ipv4 unicast
route-policy policy_2 in
!
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
if (destination in prefix_set_ex1) then
set local-preference 100
endif
if (as-path in as_path_set_ex1) then
set community (10:333) additive
endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
if (destination in prefix_set_ex1) then
if (community matches-any comm_set_ex1) then
set community (10:666) additive
endif
if (extcommunity matches-any ext_comm_set_rt_ex1) then
set community (10:999) additive
endif
endif
end-policy
!

```

Given this sample configuration, the `show rpl as-path-set as_path_set_ex1 attachpoints` command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1 attachpoints

BGP Attachpoint:Neighbor

Neighbor/Group  type  afi/safi  in/out  referring policy  attached policy
-----
10.0.101.2      --    IPv4/uni  in      policy_1          policy_1
10.0.101.3      --    IPv4/uni  in      policy_2          policy_2

```

This table describes the significant fields shown in the display.

Table 174: show rpl as-path-set attachpoints Field Descriptions

| Field | Description |
|------------------|--|
| BGP Attachpoint | Location of the attach point. |
| Neighbor/Group | IP address of the attach point on the neighbor. |
| type | Displays the address family mode. |
| afi/safi | Address family identifier or subsequent address family identifier. |
| in/out | Import or export policy. |
| referring policy | Policy that refers to the AS path set. |
| attached policy | Policy used at the attach point. |

Related Commands

| Command | Description |
|---|---|
| show rpl community-set attachpoints, on page 1723 | Displays all the policies used at an attach point that reference the named community set. |
| show rpl route-policy attachpoints, on page 1771 | Displays all the policies used at an attach point that reference the named policy. |
| show rpl prefix-set attachpoints, on page 1755 | Displays all the policies used at an attach point that reference the named prefix set. |

show rpl as-path-set references

To list all of the policies that reference the named AS path set, use the **show rpl as-path-set references** command in EXEC mode.

show rpl as-path-set *name* **references** [**brief**]

Syntax Description

| | |
|--------------|--|
| <i>name</i> | Name of the prefix set. |
| brief | (Optional) Limits the output to just the brief table and not the detailed information for the named AS path set. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl as-path-set references** command to display all policies that reference the named AS path set either directly or indirectly.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the AS path set.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
```

```

remote-as 100
address-family ipv4 unicast
  route-policy policy_1 in
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy

```

Given this sample configuration, the **show rpl as-path-set as_path_set_ex1 references** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl as-path-set as_path_set_ex1 references

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

-----
Usage/Status          count
-----
Direct                1
Indirect              0

ACTIVE                1
INACTIVE              0
UNUSED                0

-----
route-policy          usage      policy status
-----
policy_1              Direct    ACTIVE

```

This table describes the significant fields shown in the display.

Table 175: show rpl as-path-set references Field Descriptions

| Field | Description |
|--------------|--|
| Usage/Status | Displays the usage and status of all policies that reference the AS path set. Values for usage are Direct or Indirect. Values for policy status are ACTIVE, INACTIVE, or UNUSED. |
| count | Number of policies that match each usage and status option. |
| route-policy | Name of the route policies that reference the AS path set. |
| usage | Type of usage for the policy. |

| Field | Description |
|---------------|-----------------------|
| policy status | Status of the policy. |

Related Commands

| Command | Description |
|---|--|
| show rpl community-set references, on page 1726 | Lists all policies that reference the named community set. |
| show rpl route-policy references, on page 1776 | Lists all policies that reference the named policy. |
| show rpl prefix-set references, on page 1758 | Lists all policies that reference the named prefix set. |

show rpl community-set

To display the configuration of community sets, use the **show rpl community-set** command in EXEC mode.

show rpl community-set [*name*] **states** [**brief**]

Syntax Description

| | |
|---------------|--|
| name | (Optional) Name of the community set. |
| states | (Optional) Shows all unused, inactive, and active states. |
| brief | (Optional) Limits the display to a list of the names of all community sets without their configurations. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of community sets without their configurations.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
```

show rpl community-set

```

        set community (10:999) additive
    endif
endif
end-policy

```

Given this sample configuration, the **show rpl community-set comm_set_ex2** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl community-set comm_set_ex2

community-set comm_set_ex2
 65501:1,
 65501:2,
 65501:3
end-set

```

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set, on page 1713 | Displays the configuration of a named AS path set. |
| show rpl extcommunity-set, on page 1729 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1753 | Displays the configuration of a named prefix set. |
| show rpl rd-set, on page 1761 | Displays the configuration of a named RD set. |
| show rpl route-policy, on page 1768 | Displays the configuration of a named route policy. |

show rpl community-set attachpoints

To display all the policies used at an attach point that reference the named community set, use the **show rpl community-set attachpoints** command in EXEC mode.

show rpl community-set *name* **attachpoints**

| | |
|---------------------------|--------------------------------------|
| Syntax Description | <i>name</i> Name of a community set. |
|---------------------------|--------------------------------------|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl community-set attachpoints** command to display all the policies used at an attach point that reference the named community set either directly or indirectly.

The community set name is required.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.3
  remote-as 12
  address-family ipv4 unicast
   route-policy policy_2 in
 !
 !
 !
```

```

route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then <<<<<
      set community (10:999) additive
    endif
  endif
end-policy
!
```

Given this sample configuration, the **show rpl community-set attachpoints** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl community-set ext_comm_set_rt_ex1 attachpoints
BGP Attachpoint:Neighbor
Neighbor/Group  type  afi/safi  in/out  referring policy  attached policy
-----
10.0.101.3      --   IPv4/uni  in      policy_2         policy_2
```

This table describes the significant fields shown in the display.

Table 176: show rpl community-set attachpoints Field Descriptions

| Field | Description |
|------------------|--|
| BGP Attachpoint | Location of the attach point. |
| Neighbor/Group | IP address of the attach point on the neighbor. |
| type | Displays the address family mode. |
| afi/safi | Address family identifier or subsequent address family identifier. |
| in/out | Import or export policy. |
| referring policy | Policy that refers to the AS path set. |
| attached policy | Policy used at the attach point. |

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set attachpoints, on page 1715 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl prefix-set attachpoints, on page 1755 | Displays all the policies used at an attach point that reference the named prefix set. |
| show rpl rd-set attachpoints, on page 1763 | Displays all the policies used at an attach point that reference the named RD set. |

| Command | Description |
|--|--|
| show rpl route-policy attachpoints, on page 1771 | Displays all the policies used at an attach point that reference the named policy. |

show rpl community-set references

To list all the policies that reference the named community set, use the **show rpl community-set references** command in EXEC mode.

show rpl community-set *name* references [brief]

Syntax Description

| | |
|--------------|--|
| <i>name</i> | Name of a community set. |
| brief | (Optional) Limits the output to just the summary table and not the detailed information for the community set. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl community-set references** command to display all the policies that reference the named community set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the community set.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy

```

Given this sample configuration, the **show rpl extcommunity-set comm_set_ex1 references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set comm_set_ex1 references
```

```

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

```

| Usage/Status | count |
|--------------|-------|
| Direct | 1 |
| Indirect | 0 |
| ACTIVE | 1 |
| INACTIVE | 0 |
| UNUSED | 0 |

| route-policy | usage | policy status |
|--------------|--------|---------------|
| policy_2 | Direct | ACTIVE |

This table describes the significant fields shown in the display.

Table 177: show rpl community-set references Field Descriptions

| Field | Description |
|--------------|--|
| Usage/Status | Displays the usage and status of all policies that reference the community set. Values for usage are Direct or Indirect. Values for status are ACTIVE, INACTIVE, and UNUSED. |
| count | Number of policies that match each usage and status option. |
| route-policy | Name of the route policies that reference the community set. |
| usage | Type of usage for the policy. |

| Field | Description |
|---------------|-----------------------|
| policy status | Status of the policy. |

Related Commands

| Command | Description |
|--|--|
| show rpl as-path-set references, on page 1718 | Lists all policies that reference the named AS path set. |
| show rpl prefix-set references, on page 1758 | Lists all policies that reference the named prefix set. |
| show rpl rd-set references, on page 1765 | Lists all policies that reference the named RD set. |
| show rpl route-policy references, on page 1776 | Lists all policies that reference the named policy. |

show rpl extcommunity-set

To display the configuration of extended community sets, use the **show rpl extcommunity-set** command in EXEC mode.

```
show rpl extcommunity-set [name [attachpoints| references]] [cost| rt| soo] [name] [brief] [states]
```

Syntax Description

| | |
|---------------------|---|
| <i>name</i> | (Optional) Name of the community set. |
| attachpoints | (Optional) Displays all attach points for this community set. |
| references | (Optional) Displays all policies that use this community set. |
| cost | (Optional) Displays all extended community cost sets. |
| rt | (Optional) Displays all extended community RT sets. |
| soo | (Optional) Displays all extended community SoO sets. |
| brief | (Optional) Limits the display to a list of the names of all extended community sets without their configurations. |
| states | (Optional) Displays all unused, inactive, and active states. |

Command Default

If an attachpoint or reference is not specified, all configured extended community sets are displayed
If a cost, RT, or SoO sets is not specified, all configured extended community sets are displayed

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of extended community sets without their configurations.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

In the following example, the configuration of an extended community is displayed for the RT community set named `ext_comm_set_rt_ex1`:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set rt ext_comm_set_rt_ex1

ext_comm_set_rt_ex1
 1.2.3.4:34
end-set
!
```

In the following example, the configuration of an extended community is displayed with all RT set objects:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set rt

Listing for all Extended Community RT Set objects

extcommunity-set rt extrt1
 66:60001
end-set
!
extcommunity-set rt rtset1
 10:615,
 10:6150,
 15.15.15.15:15
end-set
!
extcommunity-set rt rtset3
 11:11,
 11.1.1.1:3
end-set
!
extcommunity-set rt extsool
 66:70001
end-set
!
extcommunity-set rt rtset11
 100:121,
 100:122,
 100:123,
 100:124,
 100:125,
 100:126,
 100:127,
 100:128,
 7.7.7.7:21
end-set
!
```

In the following example, the configuration of an extended community is displayed with all cost set objects:

```
RP/0/RSP0/CPU0:router# show rpl extcommunity-set cost

Listing for all Extended Community COST Set objects

extcommunity-set cost costset1
 IGP:90:914,
 Pre-Bestpath:91:915
end-set
```

```

!
extcommunity-set cost costset2
  IGP:92:916,
  Pre-Bestpath:93:917,
  IGP:94:918,
  Pre-Bestpath:95:919
end-set
!

```

In the following example, the configuration of an extended community is displayed with all SoO set objects:

Extended Community SoO Set objects

```

extcommunity-set soo sooset1
  10:151,
  100.100.100.1:153
end-set
!
extcommunity-set soo sooset3
  11:11,
  11.1.1.1:3
end-set
!

```

Related Commands

| Command | Description |
|--|--|
| show rpl as-path-set, on page 1713 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1721 | Displays the configuration of a named community set. |
| show rpl prefix-set, on page 1753 | Displays the configuration of a named prefix set. |
| show rpl rd-set, on page 1761 | Displays the configuration of a named RD set. |
| show rpl route-policy, on page 1768 | Displays the configuration of a named route policy. |

show rpl inactive as-path-set

To display the AS path sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive as-path-set** command in EXEC mode.

show rpl inactive as-path-set [detail]

Syntax Description

| | |
|---------------|---|
| detail | (Optional) Displays the content of the object and all referenced objects for inactive AS path sets. |
|---------------|---|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive as-path-set** command to display all AS path sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample1
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl inactive as-path-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl inactive as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are INACTIVE
-----
as_path_set_ex2

```

Related Commands

| Command | Description |
|--|---|
| show rpl inactive community-set, on page 1735 | Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive extcommunity-set, on page 1738 | Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive prefix-set, on page 1741 | Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive rd-set, on page 1744 | Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive route-policy, on page 1746 | Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl inactive community-set

To display the community sets that are referenced by a policy but not any policy that is used at an attach point, use the **show rpl inactive community-set** command in EXEC mode.

show rpl inactive community-set [detail]

| | |
|---------------------------|---|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for inactive community sets. |
|---------------------------|---|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive community-set** command to display all community sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
 !
 !
 neighbor 10.0.101.3
```

show rpl inactive community-set

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
route-policy sample2
  if (destination in sample2) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample2
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl inactive community-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl inactive community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are INACTIVE
-----
comm_set_ex2

```

Related Commands

| Command | Description |
|--|---|
| show rpl inactive as-path-set, on page 1732 | Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive extcommunity-set, on page 1738 | Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive prefix-set, on page 1741 | Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive rd-set, on page 1744 | Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive route-policy, on page 1746 | Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl inactive extcommunity-set

To display the extended community sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive extcommunity-set** command in EXEC mode.

show rpl inactive extcommunity-set [detail]

Syntax Description

| | |
|---------------|--|
| detail | (Optional) Displays the content of the object and all referenced objects for inactive extended community sets. |
|---------------|--|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive extcommunity-set** command to display all extended community sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
route-policy sample3
  if (destination in sample3) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample3
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl inactive extcommunity-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl inactive extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are INACTIVE
-----
ext_comm_set_rt_ex2

```

Related Commands

| Command | Description |
|---|--|
| show rpl inactive as-path-set, on page 1732 | Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive community-set, on page 1735 | Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive prefix-set, on page 1741 | Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive rd-set, on page 1744 | Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive route-policy, on page 1746 | Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl inactive prefix-set

To display the prefix sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive prefix-set** command in EXEC mode.

show rpl inactive prefix-set [detail]

| | |
|---------------------------|--|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for inactive prefix sets. |
|---------------------------|--|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive prefix-set** command to display all prefix sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
 !
 !
 neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
route-policy sample4
  if (destination in sample4) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample4
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl inactive prefix-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl inactive prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are INACTIVE
-----
sample4
prefix_set_ex2

```

Related Commands

| Command | Description |
|--|---|
| show rpl inactive as-path-set, on page 1732 | Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive community-set, on page 1735 | Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive extcommunity-set, on page 1738 | Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive route-policy, on page 1746 | Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive rd-set, on page 1744 | Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl inactive rd-set

To display the route distinguisher (RD) sets that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive rd-set** command in EXEC mode.

show rpl inactive rd-set [detail]

Syntax Description

| | |
|---------------|--|
| detail | (Optional) Displays the content of the object and all referenced objects for inactive RD sets. |
|---------------|--|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive rd-set** command to display all RD sets that are not in use at an attach point either directly or indirectly but are referenced by at least one policy in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
rd-set rdset1
 10:151,
 100.100.100.1:153,
 100.100.100.62/31:63
end-set
!
rd-set rdset2
 10:152,
 100.100.100.1:154,
 100.100.100.62/31:89
```

```
end-set
!
```

Given this sample configuration, the **show rpl inactive rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl inactive rd-set
```

```
ACTIVE -- Referenced by at least one policy which is attached INACTIVE -- Only referenced
by policies which are not attached UNUSED -- Not attached (directly or indirectly) and not
referenced
```

```
The following rd-sets are INACTIVE
```

```
-----
rdset1
rdset2
```

Related Commands

| Command | Description |
|--|---|
| show rpl inactive as-path-set, on page 1732 | Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive community-set, on page 1735 | Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive extcommunity-set, on page 1738 | Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive prefix-set, on page 1741 | Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive route-policy, on page 1746 | Displays the route policies that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl inactive route-policy

To display the route policies that are referenced by a policy but not in any policy that is used at an attach point, use the **show rpl inactive route-policy** command in EXEC mode.

show rpl inactive route-policy [detail]

Syntax Description

| | |
|---------------|---|
| detail | (Optional) Displays the content of the object and all referenced objects for inactive route policies. |
|---------------|---|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl inactive route-policy** command to display all policies that are not in use at an attach point either directly or indirectly but are referenced by at least one other policy in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
route-policy sample3
  if (destination in sample3) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if destination in prefix_set_ex1 then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample3
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl inactive route-policy** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl inactive route-policy

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following policies are (INACTIVE)
-----
sample3
policy_3

```

Related Commands

| Command | Description |
|--|---|
| show rpl inactive as-path-set, on page 1732 | Displays the AS path sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive community-set, on page 1735 | Displays the community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive extcommunity-set, on page 1738 | Displays the extended community sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive prefix-set, on page 1741 | Displays the prefix sets that are referenced by a policy but not in any policy that is used at an attach point. |
| show rpl inactive rd-set, on page 1744 | Displays the RD sets that are referenced by a policy but not in any policy that is used at an attach point. |

show rpl maximum

To display the maximum limits for lines of configuration and number of policies, use the **show rpl maximum** command in EXEC mode.

show rpl maximum [**lines**| **policies**]

| Syntax Description | |
|--------------------|---|
| lines | (Optional) Displays the number of lines of configuration limit. |
| policies | (Optional) Displays the number of policies limit. |

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl maximum** command to display the current total, current limit, and maximum limit for lines of configuration and policies.

Use the optional **lines** keyword to limit the display to the number of lines of configuration limits. Use the optional **policies** keyword to limit the display to the number of policies limits.

| Task ID | Task ID | Operations |
|---------|--------------|------------|
| | route-policy | read |

Examples The following example shows sample output from the **show rpl maximum** command:

```
RP/0/RSP0/CPU0:router# show rpl maximum
              Current      Current      Max
              Total        Limit        Limit
-----
Lines of configuration      3         65536      131072
Policies                    1          3500         5000
```

```
Compiled policies size (kB)          0
```

Table 178: [show rpl maximum Field Descriptions, on page 1750](#) describes the significant fields shown in the display.

Table 178: show rpl maximum Field Descriptions

| Field | Description |
|-----------------------------|---|
| Lines of configuration | Displays the current total, current limit, and maximum limit of lines for the policy. |
| Policies | Displays the current total, current limit, and maximum limit of policies. |
| Compiled policies size (kB) | Displays the current compiled total for policies in kilobytes. |

Related Commands

| Command | Description |
|---|---|
| rpl maximum, on page 1637 | Configures the maximum number of lines of configuration and number of policies. |

show rpl policy-global references

To display policy-global definitions, use the **show rpl policy-global references** command in EXEC mode.

show rpl policy-global references [brief]

| | |
|---------------------------|---|
| Syntax Description | brief (Optional) Limits the display to a list of the policy names. |
|---------------------------|---|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```

policy-global
  infinity '16'
end-global
!
route-policy set-rip-unreachable
  set rip-metric $infinity
end-policy
!

```

Given this sample configuration, the **show rpl policy-global references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl policy-global references
```

```
Usage Direct -- Reference occurs in this policy Usage Indirect -- Reference occurs via an
```

apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached) Status ACTIVE -- Policy is actively used at an attachpoint Status INACTIVE -- Policy is applied by an unattached policy

| Usage/Status | count |
|--------------|-------|
| ----- | ----- |
| Direct | 1 |
| Indirect | 0 |
| ACTIVE | 0 |
| INACTIVE | 0 |
| UNUSED | 1 |

| Usage | Status | Route-policy |
|--------|--------|---------------------|
| ----- | ----- | ----- |
| Direct | UNUSED | set-rip-unreachable |

show rpl prefix-set

To display the configuration of prefix sets, use the **show rpl prefix-set** command in EXEC mode.

show rpl prefix-set [*name*] **states** [**brief**]

Syntax Description

| | |
|---------------|---|
| name | (Optional) Name of the prefix set. |
| states | (Optional) Shows all unused, inactive, and active states. |
| brief | (Optional) Limits the display to a list of the names of all extended community sets without their configurations. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Because sets cannot hierarchically reference other sets or policies, no **detail** keyword exists as with the **show rpl policy** command.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

In the following example, the configuration of prefix set pset1 is displayed:

```
RP/0/RSP0/CPU0:router# show rpl prefix-set pset1
!
prefix-set pset1
 10.0.0.1/0,
 10.0.0.2/0 ge 25 le 32,
 10.0.0.5/8 ge 8 le 32,
```

show rpl prefix-set

```
10.168.0.0/16 ge 16 le 32,  
172.16.0.9/20 ge 20 le 32,  
192.168.0.5/20 ge 20 le 32  
end-set
```

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set, on page 1713 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1721 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1729 | Displays the configuration of a named extended community set. |
| show rpl route-policy, on page 1768 | Displays the configuration of a named route policy. |

show rpl prefix-set attachpoints

To display all the policies used at an attach point that reference the named prefix set, use the **show rpl prefix-set attachpoints** command in EXEC mode.

show rpl prefix-set *name* **attachpoints**

| | | |
|---------------------------|-------------|-----------------------|
| Syntax Description | <i>name</i> | Name of a prefix set. |
|---------------------------|-------------|-----------------------|

Command Default No default behavior or values

Command Modes EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl prefix-set attachpoints** command to display all the policies used at an attach point that reference the named prefix set either directly or indirectly.

The prefix set name is required.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

show rpl prefix-set attachpoints

```

remote-as 12
address-family ipv4 unicast
route-policy policy_2 in
!
!
!
route-policy policy_1
if (destination in prefix_set_ex1) then
set local-preference 100
endif
if (as-path in as_path_set_ex1) then
set community (10:333) additive
endif
end-policy
!
route-policy policy_2
if (destination in prefix_set_ex1) then
if (community matches-any comm_set_ex1) then
set community (10:666) additive
endif
if (extcommunity matches-any ext_comm_set_rt_ex1) then
set community (10:999) additive
endif
endif
end-policy

```

Given this sample configuration, the `show rpl prefix-set prefix_set_ex1 attachpoints` command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl prefix-set prefix_set_ex1 attachpoints

BGP Attachpoint:Neighbor

Neighbor/Group  type  afi/safi  in/out  referring policy  attached policy
-----
10.0.101.2      --   IPv4/uni  in      policy_1         policy_1
10.0.101.3      --   IPv4/uni  in      policy_2         policy_2

```

This table describes the significant fields shown in the display.

Table 179: show rpl prefix-set attachpoints Field Descriptions

| Field | Description |
|------------------|--|
| BGP Attachpoint | Location of the attach point. |
| Neighbor/Group | IP address of the attach point on the neighbor. |
| type | Address family mode. |
| afi/safi | Address family identifier or subsequent address family identifier. |
| in/out | Import or export policy. |
| referring policy | Policy that refers to the AS path set. |
| attached policy | Policy used at the attach point. |

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set attachpoints, on page 1715 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl community-set attachpoints, on page 1723 | Displays all the policies used at an attach point that reference the named community set. |
| show rpl route-policy attachpoints, on page 1771 | Displays all the policies used at an attach point that reference the named policy. |

show rpl prefix-set references

To list all the policies that reference the named prefix set, use the **show rpl prefix-set references** command in EXEC mode.

show rpl prefix-set *name* references [brief]

Syntax Description

| | |
|--------------|---|
| name | Name of the prefix set. |
| brief | (Optional) Limits the output to just a summary table and not the detailed information for the named prefix set. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl prefix-set references** command to list all the policies that reference the named prefix set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the named prefix set.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
 10.0.0.0/16 le 32
end-set
prefix-set too-specific
 0.0.0.0/0 ge 25 le 32
end-set
```

```

route-policy example-one
  if destination in ten-net then
    drop
  else
    set local-preference 200
    apply set-comms
  endif
end-policy
route-policy set-comms
  set community (10:1234) additive
end-policy

route-policy example-three
  if destination in too-specific then
    drop
  else
    apply example-one
    pass
  endif
end-policy

```

The following example displays information showing the usage and status of each policy that references the prefix set `ten-net`. The **brief** keyword limits the display to just a summary table and not the detailed information for the prefix set.

```
RP/0/RSP0/CPU0:router# show rpl prefix-set ten-net references brief
```

```

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

```

```

      Usage/Status          count
-----
      Direct                1
      Indirect              1

      ACTIVE                0
      INACTIVE              1
      UNUSED                1

```

This table describes the significant fields shown in the display.

Table 180: show rpl prefix-set name references Field Descriptions

| Field | Description |
|--------------|--|
| Usage/Status | Displays the usage and status of all policies that reference the prefix set. |
| count | Number of policies that match each usage and status option. |

Related Commands

| Command | Description |
|--|--|
| show rpl as-path-set references , on page 1718 | Lists all the policies that reference the named AS path set. |

| Command | Description |
|---|--|
| show rpl community-set references, on page 1726 | Lists all the policies that reference the named community set. |
| show rpl route-policy references, on page 1776 | Lists all the policies that reference the named policy. |

show rpl rd-set

To display the configuration of route distinguisher (RD) sets, use the **show rpl rd-set** command in EXEC mode.

show rpl rd-set [*name*] **states** **brief**]

Syntax Description

| | |
|---------------|---|
| name | (Optional) Name of the RD set. |
| states | (Optional) Shows all unused, inactive, and active states. |
| brief | (Optional) Limits the display to a list of the names of all RD sets without their configurations. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Because sets cannot hierarchically reference other sets or policies, no **detail** keyword exists as with the **show rpl policy** command.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

In the following example, the configuration of RD set rdset1 is displayed:

```
RP/0/RSP0/CPU0:router# show rpl rd-set rdset1
rd-set rdset1
  10:151,
  100.100.100.1:153,
```

```
show rpl rd-set
```

```
    100.100.100.62/31:63
end-set
```

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set, on page 1713 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1721 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1729 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1753 | Displays the configuration of a named prefix set. |
| show rpl route-policy, on page 1768 | Displays the configuration of a named route policy. |

show rpl rd-set attachpoints

To display all the policies used at an attach point that reference the named route distinguisher (RD) set, use the **show rpl rd-set attachpoints** command in EXEC mode.

show rpl rd-set *name* **attachpoints**

| | | |
|---------------------------|-------------|--------------------|
| Syntax Description | <i>name</i> | Name of an RD set. |
|---------------------------|-------------|--------------------|

Command Default No default behavior or values

Command Modes EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl rd-set attachpoints** command to display all the policies used at an attach point that reference the named RD set either directly or indirectly.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples This example shows the following sample configuration:

```
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
end-policy

router bgp 10
address-family vpnv4 unicast
exit
neighbor 10.0.101.1
```

```

remote-as 11
address-family vpnv4 unicast
route-policy rdsetmatch in
!

```

Given this sample configuration, the **show rpl rd-set rdset1 attachpoints** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl rd-set rdset attachpoints

BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi  in/out  vrf name
-----
10.0.101.1      --   IPv4/vpn  in      default

```

This table describes the significant fields shown in the display.

Table 181: show rpl rd-set attachpoints Field Descriptions

| Field | Description |
|----------------|---|
| Neighbor/Group | BGP neighbor or neighbor group where the specified RD is used. |
| afi/safi | BGP address family or subaddress family where the RD set is used. |
| in/out | Direction |
| vrf name | VRF name where the RD set is used. |

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set attachpoints, on page 1715 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl community-set attachpoints, on page 1723 | Displays all the policies used at an attach point that reference the named community set. |
| show rpl prefix-set attachpoints, on page 1755 | Displays all the policies used at an attach point that reference the named prefix set. |
| show rpl route-policy attachpoints, on page 1771 | Displays all the policies used at an attach point that reference the named policy. |

show rpl rd-set references

To list all the policies that reference the named route distinguisher (RD) set, use the **show rpl rd-set references** command in EXEC mode.

show rpl rd-set *name* **references** [**brief**]

Syntax Description

| | |
|--------------|---|
| name | Name of the RD set. |
| brief | (Optional) Limits the output to just a summary table and not the detailed information for the RD set. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl rd-set references** command to list all the policies that reference the named RD set.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the named RD set.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
route-policy rdsetmatch
  if rd in rdset1 then
    set community (10:112)
  elseif rd in rdset2 then
    set community (10:223)
  endif
```

```

end-policy
!
router bgp 10
  address-family vpnv4 unicast
  !
  neighbor 10.0.101.1
    remote-as 11
    address-family vpnv4 unicast
    route-policy rdsetmatch in
  !

```

Given this sample configuration, the **show rpl rd-set rdset1 references** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl rd-set rdset1 references
```

```

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

```

```

-----
Usage/Status          count
-----
Direct                1
Indirect              0

ACTIVE                1
INACTIVE              0
UNUSED                0

-----
route-policy          usage      policy status
-----
rdsetmatch            Direct    ACTIVE

```

This table describes the significant fields shown in the display.

Table 182: show rpl rd-set name references Field Descriptions

| Field | Description |
|---------------|---|
| route-policy | Name of the route policy. |
| usage | Type of reference usage for the route policy. |
| policy status | Status of the route policy. |

Related Commands

| Command | Description |
|---|--|
| show rpl as-path-set references, on page 1718 | Lists all the policies that reference the named AS path set. |
| show rpl community-set references, on page 1726 | Lists all the policies that reference the named community set. |

| Command | Description |
|--|---|
| show rpl prefix-set references, on page 1758 | Lists all the policies that reference the named prefix set. |
| show rpl route-policy references, on page 1776 | Lists all policies that reference the named policy. |

show rpl route-policy

To display the configuration of route policies, use the **show rpl route-policy** command in EXEC mode.

show rpl route-policy [*name* [**detail**]] **states** | **brief**]

Syntax Description

| | |
|---------------|---|
| name | (Optional) Name of a route policy. |
| detail | (Optional) Displays the configuration of all policies and sets that a policy uses. |
| states | (Optional) Shows all unused, inactive, and active states. |
| brief | (Optional) Limits the display to a list of the names of all extended community sets without their configurations. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the optional **brief** keyword to limit the display to a list of the names of policies without their configurations.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

In the following example, the configuration of a route policy named `policy_1` is displayed.

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1
route-policy policy_1
  if destination in prefix_set_1 and not destination in sample1 then
```

```

    if as-path in aspath_set_1 then
        set local-preference 300
        set origin igp
    elseif as-path in as_allowed then
        set local-preference 400
        set origin igp
    else
        set origin igp
    endif
else
    drop
endif
set med 120
set community (8660:612) additive
apply set_lpref_from_comm
end-policy

```

If the optional **detail** keyword is used, all routing policy language (RPL) policies and sets that route policy policy_1 uses are displayed, as shown in the following example.

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1 detail
```

```

!
prefix-set sample1
 0.0.0.0/0,
 0.0.0.0/0 ge 25 le 32,
 10.0.0.0/8 ge 8 le 32,
 192.168.0.0/16 ge 16 le 32,
 224.0.0.0/20 ge 20 le 32,
 240.0.0.0/20 ge 20 le 32
end-set
!
prefix-set prefix_set_1
 10.0.0.1/24 ge 24 le 32,
 10.0.0.5/24 ge 24 le 32,
 172.16.0.1/24 ge 24 le 32,
 172.16.5.5/24 ge 24 le 32,
 172.16.20.10/24 ge 24 le 32,
 172.30.0.1/24 ge 24 le 32,
 10.0.20.10/24 ge 24 le 32,
 172.18.0.5/24 ge 24 le 32,
 192.168.0.1/24 ge 24 le 32,
 192.168.20.10/24 ge 24 le 32,
 192.168.200.10/24 ge 24 le 32,
 192.168.255.254/24 ge 24 le 32
end-set
!
as-path-set as_allowed
 ios-regex '.*_1239_.*',
 ios-regex '.*_3561_.*',
 ios-regex '.*_701_.*',
 ios-regex '.*_666_.*',
 ios-regex '.*_1755_.*',
 ios-regex '.*_1756_.*'
end-set
!
as-path-set aspath_set_1
 ios-regex '._9148_',
 ios-regex '._5870_',
 ios-regex '._2408_',
 ios-regex '._2531_',
 ios-regex '._197_',
 ios-regex '._2992_'
end-set
!
route-policy set_lpref_from_comm
 if community matches-any (2:50) then
   set local-preference 50
 elseif community matches-any (2:60) then
   set local-preference 60
 elseif community matches-any (2:70) then

```

```

    set local-preference 70
  elseif community matches-any (2:80) then
    set local-preference 80
  elseif community matches-any (2:90) then
    set local-preference 90
  endif
end-policy
!
route-policy policy_1
  if destination in prefix_set_1 and not destination in sample1 then
    if as-path in aspath_set_1 then
      set local-preference 300
      set origin igp
    elseif as-path in as_allowed then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
  endif
  set med 120
  set community (8660:612) additive
  apply set_lpref_from_comm
end-policy

```

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set, on page 1713 | Displays the configuration of a named AS path set. |
| show rpl community-set, on page 1721 | Displays the configuration of a named community set. |
| show rpl extcommunity-set, on page 1729 | Displays the configuration of a named extended community set. |
| show rpl prefix-set, on page 1753 | Displays the configuration of a named prefix set. |

show rpl route-policy attachpoints

To display all the policies used at an attach point that reference the named policy, use the **show rpl route-policy attachpoints** command in EXEC mode.

show rpl route-policy *name* **attachpoints**

| | | |
|---------------------------|-------------|-------------------|
| Syntax Description | <i>name</i> | Name of a policy. |
|---------------------------|-------------|-------------------|

Command Default No default behavior or values

Command Modes EXEC

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy attachpoints** command to display all the policies used at an attach point that reference the named policy either directly or indirectly.

The policy name is required.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

show rpl route-policy attachpoints

```

remote-as 12
address-family ipv4 unicast
route-policy policy_2 in
!
!
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1

route-policy policy_1
if (destination in prefix_set_ex1) then
set local-preference 100
endif
if (as-path in as_path_set_ex1) then
set community (10:333) additive
endif
end-policy
!
RP/0/RSP0/CPU0:router# show rpl route-policy policy_2

route-policy policy_2
if (destination in prefix_set_ex1) then
if (community matches-any comm_set_ex1) then
set community (10:666) additive
endif
if (extcommunity matches-any ext_comm_set_rt_ex1) then
set community (10:999) additive
endif
endif
end-policy
!

```

The following command displays the route policy attach points for policy_2:

```

RP/0/RSP0/CPU0:router# show rpl route-policy policy_2 attachpoints

BGP Attachpoint: Neighbor

Neighbor/Group  type  afi/safi  in/out  vrf name
-----
10.0.101.2      --    IPv4/uni  in      default
10.0.101.2      --    IPv4/uni  out     default

```

This table describes the significant fields shown in the display.

Table 183: show rpl route-policy attachpoints Field Descriptions

| Field | Description |
|-----------------|--|
| BGP Attachpoint | Location of the attach point. |
| Neighbor/Group | IP address of the attach point on the neighbor. |
| type | Displays the address family mode. |
| afi/safi | Address family identifier or subsequent address family identifier. |
| vrf name | Name of the VPN routing and forwarding (VRF) instance. |

Related Commands

| Command | Description |
|---|---|
| show rpl as-path-set attachpoints, on page 1715 | Displays all the policies used at an attach point that reference the named AS path set. |
| show rpl community-set attachpoints, on page 1723 | Displays all the policies used at an attach point that reference the named community set. |
| show rpl prefix-set attachpoints, on page 1755 | Displays all the policies used at an attach point that reference the named prefix set. |

show rpl route-policy inline

To display all policies and sets that a policy uses expanded inline, use the **show rpl route-policy inline** command in EXEC mode.

show rpl route-policy *name* inline

Syntax Description

| | |
|-------------|-------------------|
| <i>name</i> | Name of a policy. |
|-------------|-------------------|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy inline** command to examine the configuration of a specified route policy. All policies and sets that a policy uses are gathered together and displayed expanded inline.

The policy name is required.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

The following command displays the route policy `policy_1`:

```
RP/0/RSP0/CPU0:router# show rpl route-policy policy_1
!
route-policy policy_1
  if destination in prefix_set_1 and not destination in martians then
    if as-path in aspath_set_1 then
      set local-preference 300
      set origin igp
    elseif as-path in as_allowed then
      set local-preference 400
```

```

        set origin igp
    else
        set origin igp
    endif
else
    drop
endif
set med 120
set community (8660:612) additive
apply set_lpref_from_comm
end-policy

```

The following command displays the route policy policy_1 and all the other sets or policies it refers too inline. Adding the inline keyword causes the configuration to be displayed inline for all RPL objects that the route-policy policy_1 uses.

```
RP/0/RSP0/CPU0:router#show rpl policy policy_1 inline
```

```

route-policy policy_1
  if destination in (91.5.152.0/24 ge 24 le 32, 91.220.152.0/24 ge 24 le 32, 61.106.52.0/24
  ge 24 le 32, 222.168.199.0/24
  ge 24 le 32, 93.76.114.0/24 ge 24 le 32, 41.195.116.0/24 ge 24 le 32, 35.92.152.0/24 ge
  24 le 32, 143.144.96.0/24 ge 24
  le 32, 79.218.81.0/24 ge 24 le 32, 75.213.219.0/24 ge 24 le 32, 178.220.61.0/24 ge 24 le
  32, 27.195.65.0/24 ge 24 le 32)
  and not destination in (0.0.0.0/0, 0.0.0.0/0 ge 25 le 32, 10.0.0.0/8 ge 8 le 32,
  192.168.0.0/16 ge 16 le 32, 224.0.0.0/20
  ge 20 le 32, 240.0.0.0/20 ge 20 le 32) then
    if as-path in (ios-regex '_9148_', ios-regex '_5870_', ios-regex '_2408_', ios-regex
    ' 2531 ', ios-regex ' 197 ',
    ios-regex ' 2992 ') then
      set local-preference 300
      set origin igp
    elseif as-path in
    (ios-regex '.*_1239_.*', ios-regex '.*_3561_.*', ios-regex '.*_701_.*', ios-regex
    '.*_666_.*', ios-regex '.*_1755_.*',
    ios-regex '.*_1756_.*') then
      set local-preference 400
      set origin igp
    else
      set origin igp
    endif
  else
    drop
  endif
set med 120
set community (8660:612) additive
# apply set_lpref_from_comm
if community matches-any (2:50) then
  set local-preference 50
elseif community matches-any (2:60) then
  set local-preference 60
elseif community matches-any (2:70) then
  set local-preference 70
elseif community matches-any (2:80) then
  set local-preference 80
elseif community matches-any (2:90) then
  set local-preference 90
endif
# end-apply set_lpref_from_comm
end-policy

```

show rpl route-policy references

To list all the policies that reference the named policy, use the **show rpl route-policy references** command in EXEC mode.

show rpl route-policy *name* references [brief]

Syntax Description

| | |
|--------------------|---|
| <i>name</i> | Name of a prefix set. |
| brief | (Optional) Limits the output to just a summary table and not the detailed information for the named policy. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy references** command to list all the policies that reference the named policy.

Use the optional **brief** keyword to limit the output to just a summary table and not the detailed information for the policy.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
 10.0.0.0/16 le 32
end-set
prefix-set too-specific
 0.0.0.0/0 ge 25 le 32
end-set
```

```

route-policy example-one
  if destination in ten-net then
    drop
  else
    set local-preference 200
    apply set-comms
  endif
end-policy
route-policy set-comms
  set community (10:1234) additive
end-policy
route-policy example-three
  if destination in too-specific then
    drop
  else
    apply example-one
    pass
  endif
end-policy

```

The following command displays information about the policy set-comms and how it is referenced:

```
RP/0/RSP0/CPU0:router# show rpl route-policy set-comms references
```

```

Usage Direct -- Reference occurs in this policy
Usage Indirect -- Reference occurs via an apply statement

Status UNUSED -- Policy is not in use at an attachpoint (unattached)
Status ACTIVE -- Policy is actively used at an attachpoint
Status INACTIVE -- Policy is applied by an unattached policy

```

| Usage/Status | count |
|--------------|-------|
| Direct | 1 |
| Indirect | 1 |
| ACTIVE | 0 |
| INACTIVE | 1 |
| UNUSED | 1 |

| route-policy | usage | policy status |
|---------------|----------|---------------|
| example-one | Direct | INACTIVE |
| example-three | Indirect | UNUSED |

The direct usage indicates that the route policy example-one directly applies the policy set-comms, that is, example-one has a line in the form `apply set-comms`. The usage Indirect indicates that the route policy example-three does not directly apply the route policy set-comms. However, the route policy example-three does apply the policy example-one, which in turn applies the policy set-comms, so there is an indirect reference from example-three to the route policy set-comms.

The status column indicates one of three states. A policy is active if it is in use at an attach point. In the example provided, neither example-one nor example-three is in use at an attach point, which leaves two possible states: UNUSED or INACTIVE. The route policy example-one is inactive because it has some other policy (example-three) that references it, but neither example-one nor any of the policies that reference it (example-one) are in use at an attach point. The route policy example-three has a status of unused because it is not used at an attach point and no other route policies in the system refer to it.

This table describes the significant fields shown in the display.

Table 184: show rpl route-policy references Field Descriptions

| Field | Description |
|---------------|---|
| Usage/Status | Displays the usage and status of all policies that reference the specified policy. Values for usage are Direct or Indirect. Values for status are ACTIVE, INACTIVE, and UNUSED. |
| count | Number of policies that match each usage and status option. |
| route-policy | One name for multiple policies that reference the specified policy. |
| usage | Type of usage for the policy. |
| policy status | Status of the policy. |

Related Commands

| Command | Description |
|---|--|
| show rpl as-path-set references, on page 1718 | Lists all policies that reference the named AS path set. |
| show rpl community-set references, on page 1726 | Lists all policies that reference the named community set. |
| show rpl prefix-set references, on page 1758 | Lists all policies that reference the named prefix set. |

show rpl route-policy uses

To display information about a specified named policy, use the **show rpl route-policy uses** command in EXEC mode.

show rpl route-policy *name* **uses** {**policies**|**sets**|**all**} [**direct**]

Syntax Description

| | |
|-----------------|--|
| <i>name</i> | Name of a policy. |
| policies | Generates a list of all policies that the named policy uses. |
| sets | Lists all named sets that are used by the policy. |
| all | Generates a list of both sets and policies that the named policy references. |
| <i>direct</i> | (Optional) Lists only the policies or sets used directly in the named policy block. Set or policy references that occur as a result of an apply statement are not listed. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl route-policy uses** command to display information about a specified named policy.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
prefix-set ten-net
```

```

    10.0.0.0/16 le 32
end-set
prefix-set too-specific
  0.0.0.0/0 ge 25 le 32
end-set
route-policy example-one
  if destination in ten-net then
    drop
  else
    set local-preference 200
    apply set-comms
  endif
end-policy
route-policy set-comms
  set community (10:1234) additive
end-policy
route-policy example-three
  if destination in too-specific then
    drop
  else
    apply example-one
    pass
  endif
end-policy

```

The following command lists the policies one and set-comms. It also lists the prefix sets too-specific and ten-net.

```
RP/0/RSP0/CPU0:router# show rpl route-policy example-three uses all
```

```

Policies directly and indirectly applied by this policy:
-----
    example-one set-comms

Sets referenced directly and indirectly
-----
(via applied policies) in this policy:

type prefix-set:
    ten-net too-specific

```

The sets example-one and set-comms are listed as policies that are used by the policy example-three. The policy example-one is listed because route policy example-three uses it in an **apply** statement. The policy set-comms is also listed because example-one applies it. Similarly, the prefix-set too-specific is used directly in the **if** statement in the policy example-three, and the prefix-set ten-net is used in the policy example-one. The optional **direct** keyword can be used to limit the output to just those sets and policies that are used within the example-three block itself, as shown in the following example:

```
RP/0/RSP0/CPU0:router# show rpl route-policy example-three uses all direct
```

```

Policies directly applied by this policy:
-----
    example-one

Sets used directly in this policy
-----
type prefix-set:
    too-specific

```

As can be seen in the output, the route policy set-comms and the prefix set ten-net are no longer included in the output when the **direct** keyword is used. The **direct** form of the command considers only those sets or policies used in the specified route policy and any additional policies or sets that may be used if you follow the hierarchy of **apply** statements.

This table describes the significant fields shown in the display.

Table 185: show rpl route-policy uses Field Descriptions

| Field | Description |
|-------|--|
| type | Displays the type used in the policy configuration. Values for type are prefix-set, community-set, extcommunity-set, and as-path-set. |

show rpl unused as-path-set

To display the AS path sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused as-path-set** command in EXEC mode.

show rpl unused as-path-set [detail]

Syntax Description

| | |
|---------------|---|
| detail | (Optional) Displays the content of the object and all referenced objects for unused AS path sets. |
|---------------|---|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused as-path-set** command to display all AS path sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
  !
  !
as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
!
as-path-set as_path_set_ex2
  ios-regex '^_65502_$',
  ios-regex '^_65503_$'
end-set
!
as-path-set as_path_set_ex3
  ios-regex '^_65504_$',
  ios-regex '^_65505_$'
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy

```

Given this sample configuration, the **show rpl unused as-path-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused as-path-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following as-path-sets are UNUSED
-----
as_path_set_ex3
```

Related Commands

| | |
|--|--|
| show rpl unused community-set, on page 1785 | Displays the community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1788 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1790 | Displays the prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1793 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1795 | Displays the route policies that are not referenced at all. |

show rpl unused community-set

To display the community sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused community-set** command in EXEC mode.

show rpl unused community-set [detail]

| | |
|---------------------------|---|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for unused community sets. |
|---------------------------|---|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused community-set** command to display all the community sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples This example shows the following sample configuration:

```
router bgp 2
 address-family ipv4 unicast
 !
 neighbor 10.0.101.2
  remote-as 100
  address-family ipv4 unicast
   route-policy policy_1 in
 !
 !
 neighbor 10.0.101.3
```

show rpl unused community-set

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
!
community-set comm_set_ex3
  65502:1,
  65502:2,
  65502:3
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample

```

```

    apply policy_3
end-policy

```

Given this sample configuration, the **show rpl unused community-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl unused community-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following community-sets are UNUSED
-----
comm_set_ex3

```

Related Commands

| Command | Description |
|--|--|
| show rpl unused as-path-set, on page 1782 | Displays the AS path sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1788 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1790 | Displays the prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1793 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1795 | Displays the route policies that are not referenced at all. |

show rpl unused extcommunity-set

To display the extended community sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused extcommunity-set** command in EXEC mode.

show rpl unused extcommunity-set [**cost**| **detail**| **rt**| **soo**]

Syntax Description

| | |
|---------------|--|
| cost | (Optional) Displays the unused extended-community cost objects. |
| rt | (Optional) Displays the unused extended community RT objects. |
| soo | (Optional) Displays the unused extended-community SoO objects. |
| detail | (Optional) Displays the content of the object and all referenced objects for unused extended community sets. |

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused extcommunity-set** command to display all extended community sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

The following is sample output for the **show rpl unused extcommunity-set** command:

```
RP/0/RSP0/CPU0:router:router# show rpl unused extcommunity-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following extcommunity-sets are UNUSED
-----
ext_comm_set_ex3
```

Related Commands

| Command | Description |
|---|---|
| show rpl unused as-path-set, on page 1782 | Displays the AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1785 | Displays community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1790 | Displays prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1793 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1795 | Displays the route policies that are not referenced at all. |

show rpl unused prefix-set

To display the prefix sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused prefix-set** command in EXEC mode.

show rpl unused prefix-set [detail]

Syntax Description

| | |
|---------------|--|
| detail | (Optional) Displays the content of the object and all referenced objects for unused prefix sets. |
|---------------|--|

Command Default

No default behavior or values

Command Modes

EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused prefix-set** command to display all prefix sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

Task ID

| Task ID | Operations |
|--------------|------------|
| route-policy | read |

Examples

This example shows the following sample configuration:

```
router bgp 2
  address-family ipv4 unicast
  !
  neighbor 10.0.101.2
    remote-as 100
    address-family ipv4 unicast
    route-policy policy_1 in
  !
  !
  neighbor 10.0.101.3
```

```

remote-as 12
address-family ipv4 unicast
  route-policy policy_2 in
!
!
!
prefix-set sample
  0.0.0.0/0,
  0.0.0.0/0 ge 25 le 32,
  10.0.0.0/8 ge 8 le 32,
  192.168.0.0/16 ge 16 le 32,
  224.0.0.0/20 ge 20 le 32,
  240.0.0.0/20 ge 20 le 32
end-set
!
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
prefix-set prefix_set_ex2
  220.220.220.0/24 ge 24 le 32,
  220.220.120.0/24 ge 24 le 32,
  220.220.130.0/24 ge 24 le 32
end-set
!
prefix-set prefix_set_ex3
  221.221.220.0/24 ge 24 le 32,
  221.221.120.0/24 ge 24 le 32,
  221.221.130.0/24 ge 24 le 32
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
  endif
end-policy

```

```

        endif
        if (extcommunity matches-any ext_comm_set_rt_ex2) then
            set community (10:999) additive
        endif
    endif
end-policy
!
route-policy policy_5
    apply sample
    apply policy_3
end-policy
-----
ext_comm_set_ex3

```

Given this sample configuration, the **show rpl unused prefix-set** command displays the following information:

```

RP/0/RSP0/CPU0:router# show rpl unused prefix-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following prefix-sets are UNUSED
-----
prefix_set_ex3

```

Related Commands

| Command | Description |
|--|--|
| show rpl unused as-path-set, on page 1782 | Displays AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1785 | Displays community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1788 | Displays extended community sets that are not referenced at all. |
| show rpl unused rd-set, on page 1793 | Displays the RD sets that are not referenced at all. |
| show rpl unused route-policy, on page 1795 | Displays the route policies that are not referenced at all. |

show rpl unused rd-set

To display the route distinguisher (RD) sets that are defined but not used by a policy at an attach point or referenced in a policy using an **apply** statement, use the **show rpl unused rd-set** command in EXEC mode.

show rpl unused rd-set [detail]

| | |
|---------------------------|--|
| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for unused RD sets. |
|---------------------------|--|

Command Default No default behavior or values

Command Modes EXEC

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **show rpl unused rd-set** command to display all of the RD sets that are not used in a policy at an attach point either directly or indirectly and are not referenced by any policies in the system.

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | route-policy | read |

Examples The **show rpl unused rd-set** command displays the following information:

```
RP/0/RSP0/CPU0:router# show rpl unused rd-set

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

The following rd-sets are UNUSED
-----
None found with this status.
```

Related Commands

| Command | Description |
|--|--|
| show rpl unused as-path-set, on page 1782 | Displays the AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1785 | Displays the community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1788 | Displays the extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1790 | Displays the prefix sets that are not referenced at all. |
| show rpl unused route-policy, on page 1795 | Displays the route policies that are not referenced at all. |

show rpl unused route-policy

To display the route policies that are defined but not used at an attach point or referenced using an **apply** statement, use the **show rpl unused route-policy** command in EXEC mode.

show rpl unused route-policy [detail]

| Syntax Description | detail (Optional) Displays the content of the object and all referenced objects for unused route policies. | | | | |
|---------------------------|---|---------|--------------|---------------|------------------------------|
| Command Default | No default behavior or values | | | | |
| Command Modes | EXEC | | | | |
| Command History | <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>Release 3.7.2</td> <td>This command was introduced.</td> </tr> </tbody> </table> | Release | Modification | Release 3.7.2 | This command was introduced. |
| Release | Modification | | | | |
| Release 3.7.2 | This command was introduced. | | | | |
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>Use the show rpl unused route-policy command to display route policies that are defined but not used at an attach point or referenced from another policy using an apply statement.</p> | | | | |
| Task ID | <table border="1"> <thead> <tr> <th>Task ID</th> <th>Operations</th> </tr> </thead> <tbody> <tr> <td>route-policy</td> <td>read</td> </tr> </tbody> </table> | Task ID | Operations | route-policy | read |
| Task ID | Operations | | | | |
| route-policy | read | | | | |

Examples

This example shows the following sample configuration:

```
RP/0/RSP0/CPU0:router# show run | begin prefix-set
Building configuration...
prefix-set prefix_set_ex1
  10.0.0.0/16 ge 16 le 32,
  0.0.0.0/0 ge 25 le 32,
  0.0.0.0/0
end-set
!
prefix-set prefix_set_ex2
```

show rpl unused route-policy

```

    220.220.220.0/24 ge 24 le 32,
    220.220.120.0/24 ge 24 le 32,
    220.220.130.0/24 ge 24 le 32
end-set
!
as-path-set as_path_set_ex1
  ios-regex '^_65500_$',
  ios-regex '^_65501_$'
end-set
!
as-path-set as_path_set_ex2
  ios-regex '^_65502_$',
  ios-regex '^_65503_$'
end-set
!
as-path-set as_path_set_ex3
  ios-regex '^_65504_$',
  ios-regex '^_65505_$'
end-set
!
community-set comm_set_ex1
  65500:1,
  65500:2,
  65500:3
end-set
!
community-set comm_set_ex2
  65501:1,
  65501:2,
  65501:3
end-set
!
extcommunity-set rt ext_comm_set_rt_ex1
  1.2.3.4:34
end-set
!
extcommunity-set rt ext_comm_set_rt_ex2
  2.3.4.5:36
end-set
!
route-policy sample
  if (destination in sample) then
    drop
  endif
end-policy
!
route-policy policy_1
  if (destination in prefix_set_ex1) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex1) then
    set community (10:333) additive
  endif
end-policy
!
route-policy policy_2
  if (destination in prefix_set_ex1) then
    if (community matches-any comm_set_ex1) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex1) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_3
  if (destination in prefix_set_ex2) then
    set local-preference 100
  endif
  if (as-path in as_path_set_ex2) then
    set community (10:333) additive
  endif
end-policy

```

```

end-policy
!
route-policy policy_4
  if (destination in prefix_set_ex2) then
    if (community matches-any comm_set_ex2) then
      set community (10:666) additive
    endif
    if (extcommunity rt matches-any ext_comm_set_rt_ex2) then
      set community (10:999) additive
    endif
  endif
end-policy
!
route-policy policy_5
  apply sample
  apply policy_3
end-policy
!
route ipv4 0.0.0.0/0 10.91.37.129
route ipv4 10.91.36.0/23 10.91.37.129
route ipv4 10.91.38.0/24 10.91.37.129
end

```

In the following example, route policies that are defined but not used at an attach point or referenced from another policy using an **apply** statement are displayed using the **show rpl unused route-policy** command.

```
RP/0/RSP0/CPU0:router# show rpl unused route-policy
```

```

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

```

```
The following policies are (UNUSED)
```

```

-----
policy_1
policy_2
policy_4
policy_5

```

Related Commands

| Command | Description |
|--|--|
| show rpl unused as-path-set, on page 1782 | Displays AS path sets that are not referenced at all. |
| show rpl unused community-set, on page 1785 | Displays community sets that are not referenced at all. |
| show rpl unused extcommunity-set, on page 1788 | Displays extended community sets that are not referenced at all. |
| show rpl unused prefix-set, on page 1790 | Displays prefix sets that are not referenced at all. |
| show rpl unused rd-set, on page 1793 | Displays the RD sets that are not referenced at all. |

source in

To test the source of a Border Gateway Protocol (BGP) route against the address contained in either a named or an inline prefix set, use the **source in** command in route-policy configuration mode.

source in {*prefix-set-name*| *inline-prefix-set*| *parameter*}

Syntax Description

| | |
|--------------------------|---|
| <i>prefix-set-name</i> | Name of a prefix set. |
| <i>inline-prefix-set</i> | Inline prefix set. The inline prefix set must be enclosed in parentheses. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **source in** command as a conditional expression within an **if** statement to test the source of the route against the data in either a named or an inline prefix set. A comparison that references a prefix set with zero elements in it returns false.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

The source of a BGP route is the IP peering address of the neighboring router from which the route was received.

The prefix set can contain both IPv4 and IPv6 prefix specifications.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, the source of a BGP route is tested against the data in the prefix set my-prefix-set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-A
RP/0/RSP0/CPU0:router(config-rpl)# if source in my-prefix-set then
```

In this example, the source of a BGP route is tested against the data in an inline IPv4 prefix set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-B
RP/0/RSP0/CPU0:router(config-rpl)# if source in (10.0.0.8, 10.0.0.20) then
```

In this example, the source of a route is tested against the data in an inline IPv6 prefix set:

```
RP/0/RSP0/CPU0:router(config)# route-policy policy-C
RP/0/RSP0/CPU0:router(config-rpl)# if source in (2001:0:0:1::/64, 2001:0:0:2::/64) then
```

Related Commands

| Command | Description |
|--|--|
| prefix-set, on page 1615 | Enters a prefix set configuration mode and defines a prefix set. |

suppress-route

To indicate that a given component of a BGP aggregate should be suppressed, use the **suppress-route** command in route-policy configuration mode.

suppress-route

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **suppress-route** command to indicate that a given component of an aggregate should be suppressed, that is, not advertised by BGP. See the [unsuppress-route, on page 1804](#) command for information on overriding the **suppress-route** command for individual neighbors.

The **suppress-route** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the destination is in 10.1.0.0/16, then the route is not advertised:

```
RP/0/RSP0/CPU0:router(config)# oute-policy check-aggregator
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# suppress-route

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl-if)# end-policy
```

Related Commands

| Command | Description |
|---|---|
| unsuppress-route , on page 1804 | Indicates that a given component of an aggregate should be unsuppressed |

tag

To match a specific tag value, use the **tag** command in route-policy configuration mode.

```
tag {eq|ge|le|is} {integer|parameter}
```

Syntax Description

| | |
|--------------------------|--|
| eq ge le is | Equal to; greater than or equal to; less than or equal to. |
| <i>integer</i> | Integer value. Range is 0 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **tag** command as a conditional expression within an **if** statement to match a specific tag value.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

A tag is a 32-bit integer that can be associated with a given route within the RIB.

The **eq** operator matches either a specific tag value or a parameter value. Its variants **ge** and **le** match a range of tag values that are either greater than or equal to or less than or equal to the supplied value or parameter.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the tag equals 10, then the condition returns true:

```
RP/0/RSP0RP0/CPU0:router(config-rpl)# if tag eq 10 then
```

unsuppress-route

To indicate that a given component of a BGP aggregate should be unsuppressed, use the **unsuppress-route** command in route-policy configuration mode.

unsuppress-route

Syntax Description This command has no arguments or keywords.

Command Default No default behavior or values

Command Modes Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **unsuppress-route** command to indicate that a given component of an aggregate should be unsuppressed, that is, allowed to be advertised by BGP again. This command affects routes that have been suppressed in the generation of BGP aggregates. If the request to unsuppress a route is encountered in a policy at a neighbor-out attach point, it guarantees that the routes that it affects are advertised to that neighbor even if that route was suppressed using the **suppress-route** command in a policy at the aggregation attach point.

The **unsuppress-route** command can be used as an action statement within an **if** statement. For a list of all action statements available within an **if** statement, see the **if** command.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

In the following example, if the destination is in 10.1.0.0/16, then the route is not advertised:

```
RP/0/RSP0/CPU0:router(config)# route-policy check-aggregate
RP/0/RSP0/CPU0:router(config-rpl)# if destination in (10.1.0.0/16) then
RP/0/RSP0/CPU0:router(config-rpl-if)# unsuppress-route

RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)# end-policy
```

Assuming that the policy is attached at a neighbor-out attach point, if the route 10.1.0.0/16 was suppressed in a policy at an aggregation attach point, 10.1.0.0/16 is advertised to the neighbor. Routes continue to be suppressed in advertisements to other BGP neighbors unless a specific policy is attached to unsuppress the route.

Related Commands

| Command | Description |
|---|---|
| suppress-route , on page 1800 | Indicates that a given component of a BGP aggregate should be suppressed. |

vpn-distinguisher is

To match a specific Border Gateway Protocol (BGP) VPN distinguisher, use the **vpn-distinguisher is** command in route-policy configuration mode.

vpn-distinguisher is {*number*|*parameter*}

Syntax Description

| | |
|------------------|---|
| <i>number</i> | Value assigned to a 32-bit unsigned integer. Range is from 1 to 4294967295. |
| <i>parameter</i> | Parameter name. The parameter name must be preceded with a "\$." |

Command Default

No default behavior or values

Command Modes

Route-policy configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vpn-distinguisher is** command as a conditional expression within an **if** statement to test the value of the origin attribute.

A VPN distinguisher is used in Layer 3 VPN networks for enhanced individual VPN control and to avoid route target mapping at AS boundaries in inter-AS VPN networks. Route target extended communities are removed at neighbor outbound and the VPN distinguisher value is applied on the BGP route as an extended community. When the route is received on a neighboring router in another AS, the VPN distinguisher is removed and mapped to a route target extended community.



Note

For a list of all conditional expressions available within an **if** statement, see the **if** command.

This command can be parameterized.

Task ID

| Task ID | Operations |
|--------------|-------------|
| route-policy | read, write |

Examples

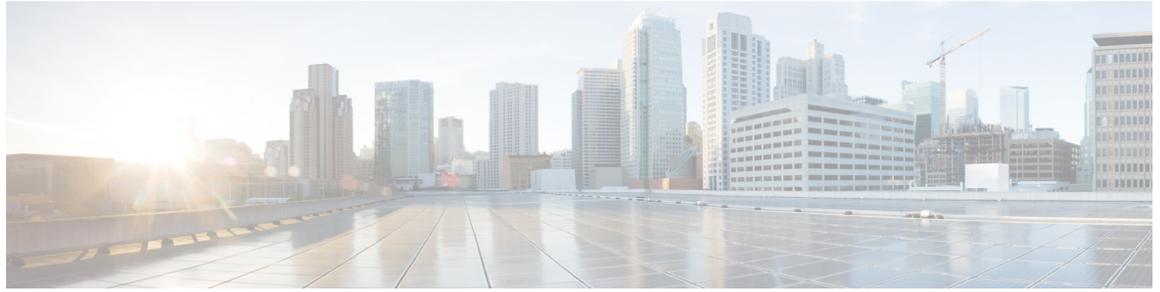
In the following example, the origin is tested within an **if** statement to learn if it is either **igp** or **egp** :

```
RP/0/RSP0/CPU0:router(config-rpl)# if origin is igp or origin is egp then
```

In the following example, a parameter is used to match a specific origin type:

```
RP/0/RSP0/CPU0:router(config)# route-policy bar($origin)
RP/0/RSP0/CPU0:router(config-rpl)# if origin is $origin then
RP/0/RSP0/CPU0:router(config-rpl-if)# set med 20
RP/0/RSP0/CPU0:router(config-rpl-if)# endif
RP/0/RSP0/CPU0:router(config-rpl)#
```

vpn-distinguisher is



Static Routing Commands

This module describes the commands used to establish static routes on Cisco ASR 9000 Series Aggregation Services Routers .

For detailed information about static routing concepts, configuration tasks, and examples, see the *Implementing Static Routes on Cisco ASR 9000 Series Router* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [address-family \(static\)](#), page 1810
- [maximum path \(static\)](#), page 1812
- [route \(static\)](#), page 1814
- [router static](#), page 1818
- [vrf \(static\)](#), page 1820

address-family (static)

To enter various address family configuration modes while configuring static routes, use the **address-family** command in the appropriate configuration mode. To disable support for an address family, use the **no** form of this command.

address-family {ipv4| ipv6} {unicast| multicast}

no address-family {ipv4| ipv6} {unicast| multicast}

Syntax Description

| | |
|------------------|---|
| ipv4 | Specifies IP Version 4 address prefixes. |
| ipv6 | Specifies IP Version 6 address prefixes. This option is available only in static router configuration mode. |
| unicast | Specifies unicast address prefixes. |
| multicast | Specifies multicast address prefixes. This option is available only in static router configuration mode. |

Command Default

All static routes belong to the default VRF if you enter address family configuration mode without entering VRF configuration mode.

Command Modes

Router static configuration
VRF router static configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **address-family** command to enter various address family configuration modes while configuring static routing sessions. From address family configuration mode, you can configure static routes using the **route** command.

Task ID

| Task ID | Operations |
|---------|-------------|
| static | read, write |

Examples

The following example shows how to enter IPv6 unicast address family mode:

```
RP/0/RSP0/CPU0:router(config)# router static  
RP/0/RSP0/CPU0:router(config-static)# address-family ipv6 unicast  
RP/0/RSP0/CPU0:router(config-static-afi)#
```

Related Commands

| Command | Description |
|---|-----------------------------|
| route (static) , on page 1814 | Establishes a static route. |

maximum path (static)

To change the maximum number of allowable static routes, use the **maximum path** command in static router configuration mode. To remove the **maximum path** command from the configuration file and restore the system to its default condition, use the **no** form of this command.

maximum path {ipv4|ipv6} *value*

no maximum path {ipv4|ipv6} *value*

Syntax Description

| | |
|--------------------|--|
| ipv4 ipv6 | Specifies IP Version 4 (IPv4) or IP Version 6 (IPv6) address prefixes. |
| <i>value</i> | Maximum number of static routes for the given AFI. The range is 1 to 140000. |

Command Default

value : 4000

Command Modes

Static router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If you use the **maximum path** command to reduce the configured maximum allowed number of static routes for a given table below the number of static routes currently configured, the change is rejected. In addition, if you commit a batch of routes that would, when grouped, push the number of static routes configured above the maximum allowed, the first *n* routes in the batch and the number previously configured are accepted, and the remainder rejected. The *n* argument is the difference between the maximum number allowed and the number previously configured.

Task ID

| Task ID | Operations |
|---------|-------------|
| static | read, write |

Examples

The following example shows how to set the maximum number of static IPv4 routes to 100000:

```
RP/0/RSP0/CPU0:router(config-static)# maximum path ipv4 100000
```

The following example shows how to remove the preceding configuration and set the maximum number of static IPv4 routes back to the default:

```
RP/0/RSP0/CPU0:router(config-static)# no maximum path ipv4 100000
```

Related Commands

| Command | Description |
|---|--|
| route (static) , on page 1814 | Enters static router configuration mode. |
| show route static | Displays the static routes in a routing table. |

route (static)

To establish static routes, use the **route** command in address family configuration mode. To remove the **route** command from the configuration, use the **no** form of this command.

```
prefix/mask [vrf vrf-name] {ip-address| type interface-path-id [ip-address| type interface-path-id] [track
track-object-name] [tunnel-id tunnel-id] [vrflabel vrf-label] [ distance ] [description text] [tag tag]
[permanent]}
```

```
no prefix/mask [vrf vrf-name] {ip-address| type interface-path-id [ip-address| type interface-path-id] [track
track-object-name] [tunnel-id tunnel-id] [vrflabel vrf-label] [ distance ] [description text] [tag tag]
[permanent]}
```

Syntax Description

| | |
|----------------------------|--|
| <i>prefix / mask</i> | <p>IP route prefix and prefix mask for the destination.</p> <p>The network mask can be specified in either of two ways:</p> <ul style="list-style-type: none"> • The network mask can be a four-part, dotted-decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means the corresponding address bit is a network address. • The network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are 1s, and the corresponding bits of the address are the network address. |
| vrf <i>vrf-name</i> | <p>(Optional) Specifies a destination VRF. This option is available when IPv4 address families are specified.</p> <p>The following names cannot be used: all, default, and global.</p> <p>The following example shows how to configure IPv4 VRF:</p> <pre>router static address-family ipv4 unicast 10.1.1.0/24 vrf vrf_a 192.168.1.1 router static vrf vrf_a address-family ipv4 unicast 10.1.1.0/24 192.168.1.1</pre> |
| <i>ip-address</i> | <p>IP address of the next hop that can be used to reach that network.</p> <ul style="list-style-type: none"> • For IPv4 address—the IP address is required, not optional, if the interface type and interface-path-id arguments are not specified. You can specify an IP address and an interface type and interface path. • For IPv6 link-local address—the interface type and interface-path-id arguments are required. The route is not valid, if the interface type and interface-path-id arguments are not specified. <p>Note A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order.</p> |
| <i>type</i> | Interface type. For more information, use the question mark (?) online help function. |

| | |
|--|---|
| <i>interface-path-id</i> | Physical interface or virtual interface. Note Use the show interfaces command to see a list of all interfaces currently configured on the router. For more information about the syntax for the router, use the question mark (?) online help function. Note A forwarding router's IP address or an interface or virtual interface path ID can be configured, in any order. |
| <i>distance</i> | (Optional) Administrative distance. Range is 1 to 254. |
| description <i>text</i> | (Optional) Specifies a description of the static route. |
| tag <i>tag</i> | (Optional) Specifies a tag value that can be used as a match for controlling redistribution using route policies. Range is 1 to 4294967295. |
| permanent | (Optional) Specifies that the route is not removed from the routing table, even if the next-hop interface shuts down or next-hop IP address is not reachable. |
| track <i>track-object-name</i> | Enables object tracking for static route. |
| tunnel-id <i>tunnel-id</i> | Specifies a Tunnel ID. |
| vrf-label <i>vrf-label</i> | Specifies a VRF label. |

Command Default

No static route is established.

vrf *vrf-name* : If no VRF is specified, the vrf where the configuration takes place is used.

Command Modes

Address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A static route is appropriate when the software cannot dynamically build a route to the destination.

Static routes have a default administrative distance of 1, in which a low number indicates a preferred route. By default, static routes are preferred to routes learned by routing protocols. You can configure an administrative distance with a static route if you want the static route to be overridden by dynamic routes. For example, you could have routes installed by the Open Shortest Path First (OSPF) protocol with an administrative distance

of 120. To have a static route that would be overridden by an OSPF dynamic route, specify an administrative distance greater than 120.

The routing table considers the static routes that point to an interface as “directly connected.” Directly connected networks are advertised by IGP routing protocols if a corresponding **interface** command is contained under the router configuration stanza of that protocol.

A static route is always associated with a VPN routing and forwarding (VRF) instance. The VRF can be the default VRF or a specified VRF. Specifying a VRF allows you to enter VRF configuration mode where you can configure a static route. If you do not specify a VRF you can configure a default VRF static route.

Use the **router static** command to configure static routes. To configure a static route, you must enter router static configuration mode and then enter an address family configuration mode or VRF configuration mode. See the **vrf (static)** command for information on configuring a static route in VRF configuration mode. After you enter an address family mode, you can enter multiple static routes. The following example shows how to configure multiple static routes in IPv4 and IPv6 address family configuration modes:

**Note**

You cannot create a VRF named default, but you can reference the default VRF.

Task ID

| Task ID | Operations |
|---------|-------------|
| static | read, write |

Examples

The following example shows how to configure IPv6 unicast address family static routes:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)# address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/2/0/7
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 GigabitEthernet0/6/0/0
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4c
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4d
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4e
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:4f
RP/0/RSP0/CPU0:router(config-static-afi)# 2b11::327a:7b00/120 2b11::2f01:50
```

Related Commands

| Command | Description |
|--|---|
| address-family (static) , on page 1810 | Enters address family configuration mode. |
| network (BGP) | Specifies a list of networks for the BGP routing process. |
| show route | Displays the current contents of the routing table. |
| show route static | Displays the static routes in a routing table. |

| Command | Description |
|--|---|
| show route summary | Displays the current contents of the routing table in summary format. |
| router static , on page 1818 | Enters router static configuration mode. |
| vrf (static) | Enters VRF static route configuration mode. |

router static

To enter static router configuration mode, use the **router static** command in global configuration mode. To remove all static route configurations and terminate the static routing process, use the **no** form of this command.

router static

no router static

Syntax Description This command has no arguments or keywords.

Command Default No static routing process is enabled.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 3.7.2 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operations |
|---------|--|-------------|
| | static | read, write |
| | bgp, ospf, hsrp, isis, vrrp, multicast, or network | read, write |

Examples The following example shows how to enter static router configuration mode:

```
RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)#
```

Related Commands

| Command | Description |
|--|---|
| address-family (static) , on page 1810 | Enters address family configuration mode. |

| Command | Description |
|---|---|
| show route | Displays the current contents of the routing table. |
| show route static | Displays the static routes in a routing table. |
| show route summary | Displays the current contents of the routing table in summary format. |
| route (static) , on page 1814 | Establishes a static route. |
| vrf (static) | Enters VRF static route configuration mode. |

vrf (static)

To configure a VPN routing and forwarding (VRF) instance and enter VRF configuration mode, use the **vrf** command in router configuration mode. To remove the VRF instance from the configuration file and restore the system to its default condition, use the **no vrf** form of this command.

vrf *vrf-name*

no vrf *vrf-name*

Syntax Description

| | |
|-----------------|---|
| <i>vrf-name</i> | Name of the VRF instance. The following names cannot be used: all, default, and global. |
|-----------------|---|

Command Default

No default behavior or values

Command Modes

Static router configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 3.7.2 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **vrf** command to configure a VRF instance. A VRF instance is a collection of VPN routing and forwarding tables maintained at the provider edge (PE) router.

A static route is always associated with a VRF, which is entirely user configurable. Static route is unique within a VRF. A static route can point to a next-hop interface, next-hop IP address, or both, which can be resided in the same VRF configured for the static route or in a different VRF. For example, routes 172.168.40.0/24 and 172.168.50.0/24 are configured as follows:

```
router static
vrf vrf_A
address_ipv4 unicast
172.168.40.0/24 loopback 1
172.168.50.0/24 vrf vrf_B 192.168.1.2
```

Routes 172.168.40.0/24 and 172.168.50.0/24 belong to vrf_A. Route 172.168.50.0/24 is not installed in vrf_A until next-hop 192.168.1.2 (a vrf_B route) is reachable.

If you are configuring a default VRF route, you do not need to enter VRF configuration mode. For example, routes 192.168.1.0/24 and 192.168.2.0/24 are configured as follows:

```
router static
```

```

address ipv4 unicast
 192.168.1.0/24 loopback 5
 192.168.2.0/24 10.1.1.1

```

Routes 192.168.1.0/24 and 192.168.2.0/24 are default VRF routes.



Note You cannot create a VRF named default, but you can reference the default VRF.

You must remove IPv4/IPv6 addresses from an interface prior to assigning, removing, or changing a VRF on an IP interface. If this is not done in advance, any attempt to change the VRF on an IP interface is rejected.

Task ID

| Task ID | Operations |
|---------|-------------|
| static | read, write |

Examples

The following example shows how to configure a VRF instance and enter VRF configuration mode:

```

RP/0/RSP0/CPU0:router(config)# router static
RP/0/RSP0/CPU0:router(config-static)# vrf vrf-1
RP/0/RSP0/CPU0:router(config-static-vrf)#

```

Related Commands

| Command | Description |
|--|--|
| address-family (static) , on page 1810 | Enters address family configuration mode and allows you to configure a static route. |



RCMD Commands

This module describes the commands used to configure and diagnose RCMD.

For detailed information about RCMD concepts, configuration tasks, and examples, see the *Implementing RCMD* module in the *Cisco ASR 9000 Series Aggregation Services Router Routing Configuration Guide*.

- [router-convergence](#), page 1824
- [monitor-convergence \(IS-IS\)](#), page 1826
- [monitor-convergence \(OSPF\)](#), page 1828
- [collect-diagnostics \(RCMD\)](#), page 1829
- [event-buffer-size \(RCMD\)](#), page 1831
- [max-events-stored \(RCMD\)](#), page 1833
- [monitoring-interval \(RCMD\)](#), page 1835
- [node disable \(RCMD\)](#), page 1837
- [prefix-list \(monitor-convergence IS-IS\)](#), page 1839
- [prefix-list \(monitor-convergence OSPF\)](#), page 1841
- [priority \(RCMD\)](#), page 1843
- [protocol \(RCMD\)](#), page 1845
- [show rcmd isis event prefix](#), page 1847
- [show rcmd ospf event prefix](#), page 1849
- [show rcmd ospf event spf](#), page 1851
- [storage-location](#), page 1854
- [track-external-routes](#), page 1856
- [track-summary-routes](#), page 1857

router-convergence

To configure route convergence monitoring and enter router convergence monitoring and diagnostics (rcmd) configuration mode, use the **router-convergence** command in global configuration mode. To remove all router convergence monitoring configurations and exit the rcmd mode, use the **no** form of this command.

router-convergence [disable]

no router-convergence

| Syntax Description | disable | [Optional] Disables the monitoring of route convergence on the entire router. |
|--------------------|---------|---|
|--------------------|---------|---|

Command Default RCMD is disabled.

Command Modes Global configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | rcmd | read, write |

Examples This example shows how to configure router-convergence command and enable rcmd configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router (config)#router-convergence
RP/0/RSP0/CPU0:router (config-rcmd)#
```

| Related Commands | Command | Description |
|------------------|--|--|
| | monitor-convergence (IS-IS) , on page 1826 | Enables route convergence monitoring for IS-IS protocol. |

| Command | Description |
|--|--|
| monitor-convergence (OSPF), on page 1828 | Enables OSPF route convergence monitoring. |

monitor-convergence (IS-IS)

To enable route convergence monitoring for IS-IS protocol, use the **monitor-convergence** command in address family configuration mode. To disable, route convergence monitoring, use the **no** form of this command.

monitor-convergence

no monitor-convergence

Syntax Description This command has no keywords or arguments.

Command Default Route convergence monitoring is disabled.

Command Modes

- Address family IPv4 unicast
- Address family IPv4 multicast
- Address family IPv6 unicast
- Address family IPv6 multicast

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | isis | read, write |

Examples This example shows how to configure route convergence monitoring for IS-IS under IPv6 multicast SAFI:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router isis isp
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv6 multicast
RP/0/RSP0/CPU0:router(config-isis-af)#monitor-convergence
```

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |
| monitor-convergence (OSPF), on page 1828 | Enables OSPF route convergence monitoring. |

monitor-convergence (OSPF)

To enable OSPF route convergence monitoring, use the **monitor-convergence** command in router OSPF configuration mode. To disable OSPF route convergence monitoring, use the **no** form of this command.

monitor-convergence

no monitor-convergence

Syntax Description This command has no keywords or arguments.

Command Default Monitor Convergence is disabled.

Command Modes Router configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | ospf | read, write |

Examples This example shows how to enable route convergence monitoring for an OSPF process:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#monitor-convergence
```

| Related Commands | Command | Description |
|------------------|--|---|
| | router-convergence , on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |
| | monitor-convergence (IS-IS) , on page 1826 | Enables route convergence monitoring for IS-IS protocol. |

collect-diagnostics (RCMD)

To collect diagnostics on specified node, use the **collect-diagnostics** command in router-convergence configuration mode. To disable collection of diagnostics, use the **no** form of this command.

collect-diagnostics *location*

no collect-diagnostics *location*

Syntax Description

| | |
|-----------------|-----------------------------------|
| <i>location</i> | Specifies the line-card location. |
|-----------------|-----------------------------------|

Command Default

Diagnostics collection is disabled.

Command Modes

Router-convergence configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

For enabling diagnostics collection on specific line-card locations, you can user can specify partially qualified semantics. However, it is not allowed to configure over-lapping locations so as to avoid errors. The following combinations of Rack and Slot are accepted:

- */*/*
- R*/*
- R/S/*

If a wildcard combination for any location is already disabled, then any other combination that overlaps with it would be rejected. For example,

- If */*/* is disabled, then all other disable commands will be rejected
- If R*/* is disabled, then disable for */*/* and R/S/* will be rejected
- If R/S/* is disabled, then disable for */*/* and R*/* will be rejected

Task ID

| Task ID | Operation |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how to enable RCMD diagnostics collection on node 0/3/CPU0:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#collect-diagnostics 0/3/CPU0
```

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

event-buffer-size (RCMD)

To specify event buffer size (in terms of number of events) for storing event traces, use the **event-buffer-size** command in router-convergence configuration mode. To disable buffer size configuration, use the **no** form of this command.

event-buffer-size *number*

no event-buffer-size

| | | |
|---------------------------|---------------|--|
| Syntax Description | <i>number</i> | Specifies the Specify the number of events. The range is 100 to 500. |
|---------------------------|---------------|--|

| | |
|------------------------|-------------|
| Command Default | 100 events. |
|------------------------|-------------|

| | |
|----------------------|----------------------------------|
| Command Modes | Router-convergence configuration |
|----------------------|----------------------------------|

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 4.2.0 | This command was introduced. |

| | |
|-------------------------|---|
| Usage Guidelines | <p>To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.</p> <p>The event-buffer-size configuration controls the ltrace buffer size. Ltraces will be stored for only the configured number of events. The default is 100 events and can be set based on the expected churn in the network. Value for event buffer impact memory usage on all RPs and monitored LCs.</p> |
|-------------------------|---|

| Task ID | Task ID | Operations |
|----------------|----------------|-------------------|
| | rcmd | read, write |

| | |
|-----------------|--|
| Examples | <p>This example shows how to configure event buffer size as 500 events:</p> <pre>RP/0/RSP0/CPU0:router#configure RP/0/RSP0/CPU0:router(config)#router-convergence RP/0/RSP0/CPU0:router(config-rcmd)#event-buffer-size 500</pre> |
|-----------------|--|

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

max-events-stored (RCMD)

To configure maximum number of events to be stored in the RCMD server, use the **max-events-stored** command in router-convergence configuration mode. To remove the number of events to be stored, use the **no** form of this command.

max-events-stored *number*

| | | |
|---------------------------|---------------|--|
| Syntax Description | <i>number</i> | Specifies the maximum number of events stored. The range is 10 to 500. |
|---------------------------|---------------|--|

| | |
|------------------------|-------------|
| Command Default | 100 events. |
|------------------------|-------------|

| | |
|----------------------|----------------------------------|
| Command Modes | Router-convergence configuration |
|----------------------|----------------------------------|

| | | |
|------------------------|----------------|------------------------------|
| Command History | Release | Modification |
| | Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The max-events-stored configuration controls the number of events that are stored in RCMD server, before the older events are deleted. The default is 100 events and can be set based on the expected churn in the network. Value for events stored impact memory usage by RCMD server.

| | | |
|----------------|----------------|-------------------|
| Task ID | Task ID | Operations |
| | rcmd | read, write |

Examples

This example shows how to configure 500 number of events to be stored in RCMD server:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#max-events-stored 500
```

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

monitoring-interval (RCMD)

To configure interval (in minutes) in which to collect logs, use the **monitoring-interval** command in router-convergence configuration mode. To disable monitoring interval configuration, use the **no** form of this command.

monitoring-interval *minutes*

no monitoring-interval *minutes*

Syntax Description

| | |
|----------------|---|
| <i>minutes</i> | Specifies the interval (in minutes) for collecting logs. The range is 5 to 120 minutes. |
|----------------|---|

Command Default

Periodic monitoring interval is 15 minutes.

Command Modes

Router-convergence configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The monitoring-interval timer controls the collection, processing, and archival (optional) of convergence data by RCMD server.

Periodic processing can get triggered if number of events detected exceed configured sizing parameters to prevent loss of data . However, this is not guaranteed since the mechanism is throttled.

To collect logs manually, use the **rcmd trigger-data-collect** command. Syslogs are generated when high churn is detected and collection mechanism is getting throttled. This indicates possible loss of data for some events. Throttling mechanism is for one processing every minute.

Task ID

| Task ID | Operations |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how to configure monitoring interval as 5 minutes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#monitoring-interval 5
```

Related Commands

| Command | Description |
|---|---|
| router-convergence , on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

node disable (RCMD)

To disable monitoring of route convergence on specified location, use the **node disable** command in router-convergence configuration mode. To reinstate, monitoring on specified location, use the **no** form of this command.

node *node-id* **disable**

no **node** *node-id* **disable**

Syntax Description

| | |
|----------------|--|
| <i>node-id</i> | Specifies line card locations for which RCMD monitoring be disabled. Disables RCMD monitoring on the specified node. No data from this node will be available in the reports that are generated. You can enter specific LCs or use wild cards. |
|----------------|--|

Command Default

Update times are gathered and reported for all LCs. Diagnostic mode is disabled on all LCs.

Command Modes

Router-convergence configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Disable monitoring on specific LCs or racks for better scalability. Disable monitoring on LCs whose update times is not going to impact the core IGP/LDP convergence that RCMD is measuring.

On LCs where monitoring is enabled, the diagnostic mode can be enabled (with threshold value) for triggering script using EEM infra for debug data collection from the router. Use diagnostic mode only for debugging purpose since it is more CPU intensive as compared to normal RCMD monitoring.

Only the following combinations of Rack and Slot are acceptable:

- */**
- R/**
- R/S/**

If a wildcard combination for any location is already disabled, then any other combination that overlaps with it would be rejected. For example,

- If */** is disabled, then all other disable commands would be rejected

- If R/*/* is disabled, then disable for */*/* and R/S/* would be rejected
- If R/S/* is disabled, then disable for */*/* and R/*/* would be rejected

Task ID

| Task ID | Operations |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how to disable monitoring on all nodes with Rack 0 and any slot (used wild card *) :

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#node 0/*/* disable
```

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

prefix-list (monitor-convergence IS-IS)

To enable individual prefix monitoring for IS-IS prefixes, use the `prefix-list` command in Router IS-IS monitor-convergence configuration mode. To disable individual prefix monitoring, use the `no` form of this command.

`prefix-list` *prefix-list-name*

`no prefix-list` *prefix-list-name*

Syntax Description

| | |
|-------------------------|--|
| <i>prefix-list-name</i> | Specifies the name of an IS-IS prefix-list. |
| Note | Configure a prefix-list under IPv4 or IPv6 using the <code>prefix-list (IP Addresses)</code> command to use for prefix monitoring. |

Command Default

All IS-IS prefixes are marked for monitoring, if the prefix-list is not configured

Command Modes

Router IS-IS monitor-convergence

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

To enable monitoring of individual prefixes, first configure a prefix-list using the `{ipv4 | ipv6} prefix-list` command. Then, use this prefix list with the `prefix-list (monitor-convergence IS-IS)`.

Task ID

| Task ID | Operation |
|---------|-------------|
| isis | read, write |

Examples

This example shows how to enable IS-IS prefix monitoring:

First, configure a prefix-list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#ipv4 prefix-list isis_monitor
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#10 permit 35.0.0.0/8 eq 32
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#commit
```

```
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#exit
```

Then, configure the prefix list command under Router IS-IS monitor-convergence configuration mode:

```
RP/0/RSP0/CPU0:router(config)#router isis isp  
RP/0/RSP0/CPU0:router(config-isis)#address-family ipv4 unicast  
RP/0/RSP0/CPU0:router(config-isis-af)#monitor-convergence  
RP/0/RSP0/CPU0:router(config-isis-af-rcmd)#prefix-list isis_monitor
```

prefix-list (monitor-convergence OSPF)

To enable individual prefix monitoring for OSPF prefixes, use the **prefix-list** command in Router OSPF monitor-convergence configuration mode. To disable individual prefix monitoring, use the **no** form of this command.

prefix-list *prefix-list-name*

no prefix-list *prefix-list-name*

Syntax Description

prefix-list-name

Specifies the name of an OSPF prefix-list.

Note Configure a prefix-list under IPv4 or IPv6 using the prefix-list (IP Addresses) command to use for prefix monitoring.

Command Default

All OSPF prefixes are marked for monitoring, if the prefix-list is not configured.

Command Modes

Router OSPF monitor-convergence

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

To enable monitoring of individual prefixes, first configure a prefix-list using the {ipv4 | ipv6} prefix-list command. Then, use this prefix list with the prefix-list (monitor-convergence OSPF).

Task ID

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This example shows how to enable OSPF prefix monitoring:

First, configure a prefix-list:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#ipv4 prefix-list ospf_monitor
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#10 permit 35.0.0.0/8 eq 32
RP/0/RSP0/CPU0:router(config-ipv4_pfx)#commit
```

```
RP/0/RSP0/CPU0:router (config-ipv4_pfx) #exit
```

Then, configure the prefix list command under Router OSPF monitor-convergence configuration mode:

```
RP/0/RSP0/CPU0:router (config) #router ospf 1  
RP/0/RSP0/CPU0:router (config-ospf) #monitor-convergence  
RP/0/RSP0/CPU0:router (config-ospf-af-rcmd) #prefix-list ospf_monitor
```

priority (RCMD)

To configure RCMD reporting parameters for low/high/critical/medium priority updates, use the **priority** command in RCMD protocol configuration mode. To disable setting up priority use the **no** form of this command.

priority {Critical| High| Low| Medium} [**disable**] [**leaf-network** *leaf-network-number*] [**threshold** *value*]
no priority {Critical| High| Low| Medium}

Syntax Description

| | |
|----------------------------|--|
| Critical | Configures the monitoring of route convergence for critical routes. |
| High | Configures the monitoring of route convergence for high priority routes. |
| Low | Configures the monitoring of route convergence for low priority routes. |
| Medium | Configures the monitoring of route convergence for medium priority routes. |
| disable | Disables the monitoring of route convergence for specified priority. |
| leaf-network | Configures the monitoring of route convergence for leaf networks. Lists up to 100 leaf networks that were added or deleted as part of SPF. |
| <i>leaf-network-number</i> | Specifies the maximum number of leaf networks monitored. The range is 10 to 100. |
| threshold | Sets the threshold value for convergence in milliseconds. If the convergence time exceeds this configured value, diagnostics collection will be triggered. |
| <i>value</i> | Specifies the threshold value (in msec). The range is 0 to 4294967295. |

Command Default

None

Command Modes

Router-convergence protocol configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **priority** command for collecting data and applying threshold for particular protocol and prefix priority. Maximum of 100 leaf networks can get logged because of scalability reasons. No default value available for threshold and this needs to be determined with deployment experience for specific network. Threshold specification is required for triggering diagnostics collection. Disable monitoring for medium and or low priority routes to help scale better. No specific order is guaranteed for leaf networks and first N prefixes that change are logged.

Task ID

| Task ID | Operations |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how Configures the monitoring of route convergence for critical routes for 100 leaf networks and at a threshold value of 1 millisecond for OSPF protocol:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#protocol OSPF
RP/0/RSP0/CPU0:router(config-rcmd-proto)#priority high
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#leaf-network 100
RP/0/RSP0/CPU0:router(config-rcmd-proto-prio)#threshold 1
```

Related Commands

| Command | Description |
|---|---|
| router-convergence , on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

protocol (RCMD)

To specify the protocol for which to configure RCMD parameters, use the **protocol** command in router-convergence configuration mode. To remove the protocol from RCMD, use the **no** form of this command.

protocol {ISIS| OSPF}

no protocol {ISIS| OSPF}

Syntax Description

| | |
|-------------|---|
| ISIS | Configures parameters related to OSPF protocol within RCMD |
| OSPF | Configures parameters related to IS-IS protocol within RCMD |

Command Default

None

Command Modes

Router-convergence configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

RCMD monitoring needs to be enabled for a specific OSPF or ISIS protocol instance.

Task ID

| Task ID | Operations |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how to enable RCMD parameters for OSPF protocol:

```
RP/0/RSP0/CPU0:router (config) #router
RP/0/RSP0/CPU0:router (config) #router-convergence
RP/0/RSP0/CPU0:router (config-rcmd) #protocol OSPF
RP/0/RSP0/CPU0:router (config-rcmd-proto) #priority high
RP/0/RSP0/CPU0:router (config-rcmd-proto-prio) #leaf-network 100
RP/0/RSP0/CPU0:router (config-rcmd-proto-prio) #threshold 1
```

Related Commands

| Command | Description |
|--|---|
| router-convergence, on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

show rcmd isis event prefix

To display the details of the individual IS-IS prefix events, use the show rcmd isis event prefix command in EXEC mode.

```
show rcmd isis isis-instance event prefix [event-number] prefix after event_number last event_number
priority {critical| high| low| medium}| threshold-exceeded] [detail] [xml]
```

Syntax Description

| | |
|---------------------------|---|
| <i>isis-instance</i> | Specifies the name of an IS-IS instance. |
| <i>event-number</i> | (Optional) Specifies the number of a specific event that is run. Range is 0-4294967295. |
| <i>prefix</i> | (Optional) Specifies events with a prefix. Specify prefix in <i>ip-address/length</i> format. |
| after | (Optional) Specifies events after a specific event number. |
| last | (Optional) Specifies the last number of events. Range is 1-500. |
| priority | (Optional) Specifies to filter events by priority. <ul style="list-style-type: none"> • critical—Event that has critical priority prefixes. • high—Event that has high priority prefixes. • low—Event that has low priority prefixes. • medium—Event that has medium priority prefixex. |
| threshold-exceeded | (Optional) Specifies events that have exceeded the threshold. |
| detail | (Optional) Provides detailed output data. |
| xml | (Optional) Provides output in XML format |

Command Default None

Command Modes EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
| isis | read |

Examples

This is sample output from the **show rcmd isis event prefix** command:

```
RP/0/RSP0/CPU0:router#show rcmd isis isp event prefix
```

show rcmd ospf event prefix

show rcmd ospf *ospf-instance* **event prefix** [*event_number*|*prefix*| **after** *event_number*| **last** *event_number*|
priority {**critical**| **high**| **low**| **medium**}| **threshold-exceeded**] [**detail**] [**xml**]

| Syntax Description | |
|---------------------------|---|
| <i>ospf-instance</i> | Specifies the name of an OSPF instance. |
| <i>event-number</i> | (Optional) Specifies the number of a specific event that is run. Range is 0-4294967295. |
| <i>prefix</i> | (Optional) Specifies events with a prefix. Specify prefix in <i>ip-address/length</i> format. |
| after | (Optional) Specifies events after a specific event number. |
| last | (Optional) Specifies the last number of events. Range is 1-500. |
| priority | (Optional) Specifies to filter events by priority. <ul style="list-style-type: none"> • critical—Event that has critical priority prefixes. • high—Event that has high priority prefixes. • low—Event that has low priority prefixes. • medium—Event that has medium priority prefixex. |
| threshold-exceeded | (Optional) Specifies events that have exceeded the threshold. |
| detail | (Optional) Provides detailed output data. |
| xml | (Optional) Provides output in XML format |

Command Default None

Command Modes EXEC

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.0 | This command was introduced. |

Usage Guidelines**Task ID**

| Task ID | Operation |
|---------|-------------|
| ospf | read, write |

Examples

This is sample output from **show rcmd ospf event prefix** command:

```

OSPF process: 1
Event: 1

Prefix: 255.255.255.255/32      Cost: 10          Priority: High
SPF Event No: 0                Route-Type: Intra Change-Type: Add

Nexthop: GigabitEthernet-2/0/0/1 Neighbor: 2.2.2.2 Change-Type: Add
GigabitEthernet-2/0/0/2 Neighbor: 1.1.1.1 Change-Type: Delete

Start time: Jan 1 05:32:22.118
Timeline:
IP Route Program Time:      Min: 40 (0/2/CPU0)   Max: 66 (0/1/CPU0)
MPLS Label Program Time:   Min: 173 (0/1/CPU0) Max: 197 (0/3/CPU0)
Details:
  RIBv4-Enter                6      <offset from Start time>
  RIBv4-Exit                 12
  RIBv4-Redist                8
  LDP Enter                  10
  LDP Exit                   16
  LSD Enter                   27
  LSD Exit                    42
  LC Details(IP Path):
    S 0/1/CPU0                66
    F 0/2/CPU0                40
    0/3/CPU0                  56
  LC Details(MPLS Path):
    F 0/1/CPU0                173
    0/2/CPU0                  174
    S 0/3/CPU0                197

```

show rcmd ospf event spf

To display route convergence monitoring and diagnostics information for OSPF shortest path first events, use the **show rcmd ospf event spf** command in EXEC mode.

show rcmd ospf *ospf-instance* **event spf** [*spf-run*] **after** | **last** | **no-route-change** | **pending** | **route-change** | **threshold-exceeded**] [**detail**] [**xml**]

Syntax Description

| | |
|-------------------------|---|
| <i>ospf-instance</i> | Specifies the OSPF instance number. |
| <i>spf-run</i> | (Optional) Specifies a specific OSPF SPF run. Range is 0-4294967295. |
| after | (Optional) Specifies events after a specific number of events. Range is 0-4294967295. |
| last | (Optional) Specifies the last "N" events. Range for "N" is 1-500. |
| no-route-change | (Optional) Displays information about events that have no-route-changes. |
| pending | (Optional) Displays events that are pending for post processing. |
| route-change | (Optional) Displays events that have route-change. |
| threshold-exceed | (Optional) Displays that have exceeded the threshold. |
| detail | (Optional) Displays detailed information about the SPF event. |
| xml | (Optional) Displays information in XML format. |

Command Default None

Command Modes EXEC

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-----------|
|---------|-----------|

Examples

This is sample output from the show rcmd ospf event spf command:

```
RP/0/RSP0/CPU0:router#show rcmd ospf 1 event spf last 1 detail

Event Status:
^ no route change # threshold exceeded ~ incomplete data * collection pending

OSPF process: 1

SPF run: 16
Trigger: Apr 10 23:07:01.614 Start: 0 Duration: 2
Dijkstra Statistics: Runs: 1 LSA changes: 0
IA/Ext Statistics: Runs: 4 LSA processed: 36
Timeline Summary:
  Priority: Critical
  Route Count: Added: 18 Deleted: 0 Modified:
0 FRR Coverage: Routes: 9/9(100%) Paths: 18/18(100%)
  IP Route Program Time: Min: 9(0/2/CPU0) Max: 11(0/3/CPU0)
  MPLS Label Program Time: Min: 18(0/1/CPU0) Max: 22(0/3/CPU0)
  Priority: High
  Route Count: Added: 18 Deleted: 0 Modified:
0 FRR Coverage: Routes: 9/9(100%) Paths: 18/18(100%)
  IP Route Program Time: Min: 11(0/1/CPU0) Max: 12(0/2/CPU0)
  MPLS Label Program Time: Min: 21(0/2/CPU0) Max: 25(0/3/CPU0)
  Priority: Medium
  Route Count: Added: 18 Deleted: 0 Modified:
0 FRR Coverage: Routes: 9/9(100%) Paths: 18/18(100%)
  IP Route Program Time: Min: 12(0/3/CPU0) Max: 15(0/2/CPU0)
  MPLS Label Program Time: Min: 22(0/2/CPU0) Max: 26(0/3/CPU0)
  Priority: Low
  Route Count: Added: 21 Deleted: 0 Modified:
```

| | | |
|--------------------------|----------------------|---------------------|
| FRR Coverage: | Routes: 10/10 (100%) | Paths: 21/21 (100%) |
| IP Route Program Time: | Min: 14 (0/1/CPU0) | Max: 19 (0/3/CPU0) |
| MPLS Label Program Time: | Min: 28 (0/1/CPU0) | Max: 33 (0/2/CPU0) |

Dijkstra Info:

| | | | |
|------------------------------|---------|-----------|-------------|
| Area: 0.0.0.0 | Run: 9 | | |
| Trigger: Apr 10 23:07:01.562 | Wait: 0 | Start: 52 | Duration: 0 |

storage-location

To specify where to store the extended routing-diagnostics that are collected when threshold exceeds, use the **storage-location** command in router-convergence configuration mode. To disable storing routing-diagnostics to a specific location, use the **no** form of this command.

storage-location [**diagnostics** *directory-path* | **diagnostics-size** *maximum-directory-size* | **reports** *directory-path* | **reports-size** *maximum-directory-size*]

no storage-location

Syntax Description

| | |
|-------------------------------|--|
| diagnostics | Specifies an absolute directory path for storing diagnostic reports. |
| <i>directory-path</i> | Specifies the path of the absolute directory for storing diagnostic reports. |
| diagnostics-size | Specifies the maximum size of diagnostics directory. |
| <i>maximum-directory-size</i> | Specified the size of the diagnostics directory. The range is 5% to 80%. |
| reports | Specifies an absolute directory path for storing reports. |
| <i>directory-path</i> | Specifies the path of the absolute directory for storing reports. |
| reports-size | Specifies the maximum size of the reports directory. The range is 5% to 80%. |

Command Default

No default storage location. Mechanism is disabled.

Command Modes

Router-convergence configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.2.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The storage location can be local disk or remote tftp space.

RCMD server can periodically archive reports in XML format for persistency. This mechanism is enabled when archival location is configured. Debug data collected in diagnostics mode is dumped to the configured diagnostic location (else it would get lost). When using local disk, the percentage of disk space to be used can

be specified, and RCMD server will delete older reports on reaching the limit. Archival (specifically on local disk) is CPU intensive. Use a remote XML server to periodically collect reports from the router and archive on the server's local storage.

Task ID

| Task ID | Operations |
|---------|-------------|
| rcmd | read, write |

Examples

This example shows how to configure storage location as *tftp://202.153.144.25/auto/tftp-chanvija-blr/rcmd/dump/reports* for reports and */hddisk:/rcmd_logs* for diagnostics:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router-convergence
RP/0/RSP0/CPU0:router(config-rcmd)#storage-location
RP/0/RSP0/CPU0:router(config-rcmd-store)#diagnostics /hddisk:/rcmd_logs
RP/0/RSP0/CPU0:router(config-rcmd-store)#reports
tftp://202.153.144.25/auto/tftp-chanvija-blr/rcmd/dump/reports
```

Related Commands

| Command | Description |
|---|---|
| router-convergence , on page 1824 | Configures route convergence monitoring and enters router convergence monitoring and diagnostics (rcmd) configuration mode. |

track-external-routes

To enable tracking of external (Type-3/5/7) LSAs prefix monitoring, use the track-external-routes command in Router OSPF monitor-convergence configuration mode. To disable, tracking of external LSAs prefix monitoring, use the no form of this command.

track-external-routes

no track-external-routes

This command has no keywords or arguments.

Command Default Route OSPF monitor-convergence

Command Modes External LSAs prefix monitoring is disabled.

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | ospf | read, write |

Examples This example shows how to enable tracking of external LSAs prefix monitoring:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#monitor-convergence
RP/0/RSP0/CPU0:router(config-ospf-af-rcmd)#track-external-routes
```

track-summary-routes

To enable tracking of summary (inter-area) routes prefix monitoring, use the track-summary-routes command in Router OSPF monitor-convergence configuration mode. To disable tracking of summary router prefix monitoring, use the no form of this command.

track-summary-routes

no track-summary-routes

This command has no keywords or arguments.

Command Default Router OSPF monitor-convergence

Command Modes Summary routes prefix monitoring is disabled.

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | ospf | read, write |

Examples This example shows how to enable tracking of summary routes prefix monitoring:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router ospf 100
RP/0/RSP0/CPU0:router(config-ospf)#monitor-convergence
RP/0/RSP0/CPU0:router(config-ospf-af-rcmd)#track-summary-routes
```




Locator/ID Separation Protocol Commands

This module describes the commands used to configure and monitor Locator/ID Separation Protocol (LISP) on Cisco IOS XR.

- [Locator/ID Separation Protocol on Cisco IOS XR](#), page 1860
- [address-family \(LISP\)](#), page 1861
- [eid-mtu](#), page 1863
- [eid-table](#), page 1865
- [etr](#), page 1867
- [etr accept-map-request-mapping](#), page 1869
- [etr map-cache-ttl](#), page 1871
- [etr map-server](#), page 1873
- [itr map-resolver](#), page 1875
- [locator reachability](#), page 1877
- [locator-set](#), page 1879
- [locator-table](#), page 1881
- [loc-reach-algorithm rloc-probing](#), page 1883
- [map-cache-limit](#), page 1885
- [map-cache](#), page 1887
- [map-request-source](#), page 1889
- [other-xtr-probe](#), page 1891
- [proxy-etr](#), page 1893
- [proxy-itr](#), page 1895
- [remote-rloc-probe](#), page 1897
- [router lisp](#), page 1899
- [solicit-map-request](#), page 1901

- [use-petr](#), page 1903

Locator/ID Separation Protocol on Cisco IOS XR

Locator/ID Separation Protocol (LISP) is a simple, incremental, network-based protocol designed to implement separation of Internet addresses into Endpoint Identifiers (EIDs) and Routing Locators (RLOCs).

LISP stands for Locator/ID Separation Protocol and is a next-generation IP routing feature that creates a new paradigm in how IP addressing is assigned and interpreted by splitting the device identity, known as an endpoint identifier (EID), and its location, known as its routing locator (RLOC), into two different namespaces. Creating separate IP addresses for EID and RLOC functions yields several advantages, including improved scalability of the routing system through greater aggregation of RLOCs and improved multihoming efficiency and ingress traffic engineering. Hosts do not have to change IP addresses and therefore, no IP address numbering costs are involved with the LISP implementation.

LISP sites use IP addresses in the EID namespace to address hosts and in Domain Name System (DNS) in exactly the same way they are currently used. These addresses are not advertised within the non-LISP RLOC namespace (that is, the Internet), but instead are advertised by the LISP mapping services. The LISP site router supports the LISP functionality of Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR).

LISP is a pull model analogous to DNS and is massively scalable. LISP is address family agnostic and can be deployed incrementally.

LISP creates a Level of indirection with two namespaces: EID and RLOC. The EID (Endpoint Identifier) is the IP address of a host. The RLOC (Routing Locator) is the IP address of the LISP router for the host. EID-to-RLOC mapping is the distributed architecture that maps EIDs to RLOCs. The LISP Map Lookup is analogous to a DNS lookup. DNS resolves IP addresses for URLs. LISP resolves locators for queried identifiers or EID prefix.

LISP in Cisco IOS XR supports:

- Proxy Ingress Tunnel Router (PITR) and Proxy Egress Tunnel Router (PETR). PITR must be configured using map resolver (no ALT support).
- Default table support for EID and RLOC space.
- The **router lisp** command in global configuration mode enables LISP configuration mode.



Note

The LISP command line interface, show commands output, and schema is to be changed in Cisco IOS XR Release 4.3.1 to be similar to the LISP command line interface on Cisco IOS.

address-family (LISP)

To enter Locator ID and separation protocol (LISP) address family configuration mode, use the **address-family** command in LISP configuration mode. To exit the LISP address family configuration mode, use the **no** form of this command.

```
address-family {ipv4| ipv6} unicast
no address-family {ipv4| ipv6} unicast
```

Syntax Description

| | |
|----------------|-----------------------------------|
| ipv4 | Selects IPv4 address family. |
| ipv6 | Select IPv6 address family. |
| unicast | Selects unicast address prefixes. |

Command Default

LISP address family configuration is disabled.

Command Modes

LISP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to enable IPv6 address family configuration for LISP:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#
```

Related Commands

| Command | Description |
|--|---|
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

eid-mtu

To configure MTU sizes for IPv4 or IPv6 LISP payloads, use the **eid-mtu** command in the LISP address family configuration mode. To remove the configured the eid-mtu size, use the **no** form of this command.

```
eid-mtu {ipv4-locator ipv4_bytes | ipv6-locator ipv6_bytes}
```

```
no eid-mtu
```

Syntax Description

| | |
|---------------------|--|
| ipv4-locator | Specifies the MTU for packets transported through IPv4 RLOC space. |
| <i>ipv4_bytes</i> | Specifies the MTU in bytes for packets transported through IPv4 RLOC space. The value ranges from 68 to 65535. |
| ipv6-locator | Specifies the MTU for packets transported through IPv6 RLOC space. |
| <i>ipv6_bytes</i> | Specifies the MTU in bytes for packets transported through IPv6 RLOC space. The value ranges from 1280 to 65535. |

Command Default

None

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure MTU size of 1300 bytes for the IPv6 LISP payloads:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#eid-mtu ipv4-locator 1320 ipv6-locator 1300
```

Related Commands

| Command | Description |
|--|---|
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

eid-table

To configure a Locator ID Separation Protocol (LISP) instance ID for association with a virtual routing and forwarding (VRF) table or default table through which the endpoint identifier (EID) address space is reachable, use the **eid-table** command in the LISP configuration mode. To remove this association, use the **no** form of this command.

```
eid-table {default| [vrf vrf_name]} instance-id instance_id
```

```
no eid-table {default| vrf vrf_name} instance-id instance_id
```

Syntax Description

| | |
|--------------------|---|
| default | Selects the default (global) routing table for association with the configured instance ID. |
| vrf | Selects the specified VRF table for association with the configured instance ID. |
| <i>vrf_name</i> | Specifies the name of the VRF. |
| instance | Specifies the instance ID to be associated with this EID table. |
| <i>instance_id</i> | Specifies the instance ID value. This value ranges between 0 and 16777215. |

Command Default

A router configured for LISP associates the default table with instance ID 0.

Command Modes

LISP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

The **eid-table** command is used to associate a LISP instance ID with either the default routing table, or a VRF table through which its EID address space is reachable. When a LISP instance ID is specified, LISP Map Registration (control plane) messages include this instance ID along with the associated EID prefixes upon registering and LISP data plane packets include this instance ID in the LISP header.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure LISP instance ID for association with a virtual routing and forwarding (VRF) table:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#eid-table vrf vrf1 instance-id 45
```

Related Commands

| Command | Description |
|--|---|
| locator-table, on page 1881 | Specifies the RLOC table. |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |

etr

To enable Egress Tunnel Router functionality, use the **etr** command in the LISP address family configuration mode. To disable the ETR functionality, use the **no** form of this command.

etr

no etr

Syntax Description This command has no keywords or arguments.

Command Default ETR functionality is not enabled by default.

Command Modes LISP IPv4 address family
LISP IPv6 address family

| Command History | Release | Modification |
|------------------------|----------------|------------------------------|
| | Release 4.3.1 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|----------------|----------------|------------------|
| | lisp | read, write |

Examples The example shows how to turn on ETR functionality in all eid-tables, unless it is explicitly disabled.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr
```

| Related Commands | Command | Description |
|-------------------------|--|--|
| | router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

| Command | Description |
|---|---|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

etr accept-map-request-mapping

To configure an Egress Tunnel Router (ETR) to cache to cache EID-to-RLOC mapping data contained, that ITRs may attach to a map-request message, use the **etr accept-map-request-mapping** command in the LISP address family configuration mode. To remove this functionality, use the **no** form of this command.

etr accept-map-request-mapping [verify]

no etr accept-map-request-mapping

Syntax Description

| | |
|---------------|--|
| verify | (Optional) Specifies that mapping data should be cached but not used for forwarding packets until the ETR can send its own map request to one of the locators from the mapping data record and receive a map reply with the same data in response. |
|---------------|--|

Command Default

No caching of mapping data in a map-request message.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

If an ETR receives a map-request message that contains mapping data for the invoking IPv6 source-EID's packet, then the ETR, by default, ignores the mapping data. However, if you configure the **etr accept-map-request-mapping** command, the ETR caches the mapping data in its map cache and immediately uses it for forwarding packets.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure **etr accept-map-request-mapping** command in the LISP IPv6 address family configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr accept-map-request-mapping verify
```

Related Commands

| Command | Description |
|--|---|
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| etr , on page 1867 | Enables Egress Tunnel Router (ETR) functionality. |

etr map-cache-ttl

To configure the time-to-live (TTL) value inserted into LISP IPv4 or IPv6 map-reply messages, use the **etr map-cache-ttl** command in the LISP address family configuration mode. To remove the configured TTL value and return to the default value, use the **no** form of this command.

etr map-cache-ttl *minutes*

no etr map-cache-ttl *minutes*

Syntax Description

| | |
|----------------|--|
| <i>minutes</i> | Specifies the value, in minutes, to be inserted in the TTL field in map-reply messages. The value ranges from 60 to 10080 minutes. |
|----------------|--|

Command Default

The default TTL value is 1440 minutes (24 hours).

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use this command to change the default value associated with the Time-to-Live (TTL) field in map-reply messages. Entering this command changes the default TTL that remote ITRs will cache and use for your site's endpoint identifier (EID) prefix. The default value is 1440 minutes (24 hours), and the minimum value is 60 minutes.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure **etr map-cache-ttl** command with TTL value of 60 minutes:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
```

```
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast  
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-cache-ttl 60
```

Related Commands

| Command | Description |
|---|---|
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| etr, on page 1867 | Enables Egress Tunnel Router (ETR) functionality. |

etr map-server

To configure the options related to the etr map-server (MS) such as locator, authentication key and whether or not the map server is allowed to respond on the ETR's behalf to map-requests (proxy-reply option), use the **etr map-server** command in the LISP address family configuration mode. To remove the configured options, use the **no** form of this command.

```
etr map-server IP {key [clear| encrypted] LINE| proxy-reply}
no etr map-server
```

Syntax Description

| | |
|--------------------|---|
| <i>ip_address</i> | Specifies the IPv4 or IPv6 address of the map server. |
| key | Specifies that an authentication key will follow either encrypted or unencrypted option. |
| LINE | Specifies the key (either encrypted or unencrypted). |
| clear | Specifies the UNENCRYPTED (cleartext) key. |
| encrypted | Specifies that an ENCRYPTED key will follow. |
| proxy-reply | Specifies that the map-server with the specified RLOC is allowed to respond to map-requests on behalf of the ETR. |

Command Default

LISP map server locator addresses are not configured by default.

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example configures 2001:0db8::1 as map-server, specifying that *supers3cr3tpassw0rd* will be used as the shared secret for HMAC calculations, and that this map-server may respond to map-requests on behalf of the ETR.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-server 2001:0db8::1 key clear
supers3cr3tpassw0rd
RP/0/RSP0/CPU0:router(config-lisp-af)#etr map-server 2001:0db8::1 proxy-reply
```

Related Commands

| Command | Description |
|--|---|
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| etr , on page 1867 | Enables Egress Tunnel Router (ETR) functionality. |

itr map-resolver

To configure an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR, when sending Map-Requests for IPv4 EID-to-RLOC mapping resolution, use the **itr map-resolver** command in LISP address family configuration mode. To remove the configured locator address of the LISP Map-Resolver, use the **no** form of this command.

itr map-resolver *map-resolver-address*

no itr map-resolver *map-resolver-address*

Syntax Description

| | |
|-----------------------------|---|
| <i>map-resolver-address</i> | Specifies the IPv4 or IPv6 locator addresses of the Map-Resolver. |
|-----------------------------|---|

Command Default

No LISP Map-Resolver locator address is configured.

Command Modes

LISP IPv4 address family configuration
LISP IPv6 address family configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example configures an ITR to use the Map-Resolver located at 10.2.3.4 when sending its Map- Request messages:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#itr map-resolver 10.2.3.4
```

This example configures and ITR to use the Map-Resolver located at 2001:DB8:0A::1 when sending its Map-Request messages:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router (config)#router lisp
RP/0/RSP0/CPU0:router (config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router (config-lisp-afi)#itr map-resolver 2001:DB8:0A::1
```

Related Commands

| Command | Description |
|--|---|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

locator reachability

To configure the conditions that must be met for a routing locator (RLOC) to be treated as being reachable, use the **locator reachability** command in the LISP address family configuration mode. To remove the configuration, use the **no** form of the command.

locator reachability exclude-default

no locator reachability exclude-default

Syntax Description

| | |
|------------------------|---|
| exclude-default | If a remote RLOC is reached via the default route, then it is treated as unreachable. |
|------------------------|---|

Command Default

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure the **locator reachability** command:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#locator reachability exclude-default
```

Related Commands

| Command | Description |
|---|--|
| loc-reach-algorithm rloc-probing , on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

locator-set

To configure a named locator set, use the **locator-set** command in the in LISP configuration mode. To disable a the locator-set, use the **no** form of the command.

```
locator-set name {ip_address {priority priority_value| weight weight_value}| auto-discover-rlocs}
no locator-set name
```

Syntax Description

| | |
|----------------------------|--|
| <i>name</i> | Specifies the name of the locator set. |
| ip_address | Specifies the RLOC IP address of Loopback or other Egress Tunnel Router (ETR) interfaces. |
| priority | Configures the preferred locators from the locator set. When multiple locators have the same priority traffic may be load-balanced across them. |
| <i>priority_value</i> | Specifies the value of the priority assigned to the RLOC. The value ranges from 0 to 255. A lower value indicates a higher priority. |
| weight | Specifies how to determine to load-share traffic between multiple locators when the priorities assigned to multiple locators are the same. |
| <i>weight_value</i> | Specifies the value of the percentage of traffic to be load-shared. The value ranges from 0 to 100. |
| auto-discover-rlocs | Configures the Egress Tunnel Router (ETR) to discover the locators of all routers configured to function as both an ETR and an Ingress Tunnel Router (ITR)--such routers are referred to as xTRs--in ETRs LISP site when the site uses multiple xTRs and each xTR is configured to use DHCP-learned locators or configured with only its own locators. |

Command Default

Command Modes

LISP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

A locator set is a named group of Routing LOCators (RLOCs). It can be used in conjunction with the **database-mapping** and **map-cache** configuration commands.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure the locator-set priority such that 2001:1::2 would have the highest priority and receive all the traffic. 2001:2::3 and 64.10.10.5 would each receive 50% of the traffic if 2001:1::2 is unreachable.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#locator-set loc1
RP/0/RSP0/CPU0:router(config-lisp)#2001:1::2 priority 1 weight 100
RP/0/RSP0/CPU0:router(config-lisp)#2001:2::3 priority 2 weight 50
RP/0/RSP0/CPU0:router(config-lisp)#64.10.10.5 priority 2 weight 50
```

Related Commands

| Command | Description |
|--|--|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

locator-table

To associate a virtual routing and forwarding (VRF) table through which the routing locator address space is reachable to a router Locator ID Separation Protocol (LISP) instantiation, use the **locator-table** command in LISP configuration mode. To remove this association, use the **no** form of this command.

locator-table *name* [**default**| **vrf** *vrf_name*]

no locator-table *name*

Syntax Description

| | |
|-----------------|--|
| default | Selects the default (global) routing table for association with the routing locator address space. |
| vrf | Selects the routing table for the specified VRF name for association with the routing locator address space. |
| <i>vrf_name</i> | Specifies the name of the VRF. |

Command Default

None

Command Modes

LISP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.



Note

The locator-table configuration is mandatory for LISP to function.

When a LISP device is deployed in a multitenant (virtualized) network environment with segmented routing locator (RLOC) address space, separate router LISP instantiations are required for each locator address space. Separate instantiations are created by including the optional id entry with the **router lisp** command. Each router LISP instantiation is considered to be standalone and must be associated with an RLOC address space. The **locator-table** command is used to associate a VRF table through which the routing locator address space is reachable to a router LISP instantiation. All necessary LISP components used in the operation of that particular router LISP instantiation, (for example, map server, map resolver, proxy ingress tunnel router (PITR), proxy egress tunnel router (PETR), and other routers that function as both egress and ingress tunnel

routers, also known as xTRs) must be reachable via the routing locator address space referred to by the **locator-table** command.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to associate a VRF table to a LISP instantiation:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#locator-table mplsvpn
```

Related Commands

| Command | Description |
|---|---|
| eid-table | Configures a Locator ID Separation Protocol (LISP) instance ID for association with a virtual routing and forwarding (VRF) table or default table through which the endpoint identifier (EID) address space is reachable. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| loc-reach-algorithm rloc-probing , on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

loc-reach-algorithm rloc-probing

To configure locator reachability algorithm, RLOC Probing, and to determine reachability status for other iBGP peers through the IGP domain, use the **loc-reach-algorithm rloc-probing** command in LISP configuration mode. To disable the locator reachability algorithm, use the **no** form of this command .

loc-reach-algorithm rloc-probing

no loc-reach-algorithm rloc-probing

Syntax Description This command has no keywords or arguments.

Command Default The locator reachability algorithm rloc-probing is disabled.

Command Modes LISP configuration

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | lisp | read, write |

Examples This example shows how to configure the location reachability algorithm RLOC probing:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#loc-reach-algorithm rloc-probing
```

| Related Commands | Command | Description |
|------------------|--|---|
| | address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

| Command | Description |
|--|---|
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

map-cache-limit

To configure the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router, use the **map-cache-limit** command in LISP address family configuration mode. To remove the configured map-cache limit, use the **no** form of this command.

map-cache-limit *map-cache-size*

no map-cache-limit *map-cache-size*

Syntax Description

| | |
|-----------------------|--|
| <i>map-cache-size</i> | Specifies the map cache size value. Range is 1 to 65535. |
|-----------------------|--|

Command Default

Map cache size: 1000

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example configures a lisp cache-limit of 2000 entries:

```
RP/0/RSP0/CPU0:router(config-lisp-afi)#map-cache-limit 2000
```

Related Commands

| Command | Description |
|--|---|
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

| Command | Description |
|--|---|
| loc-reach-algorithm rloc-probing , on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit , on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache , on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr , on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr , on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

map-cache

To configure a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or to statically configure the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix, use the **map-cache** command in LISP address family configuration mode. To remove the configuration, use the **no** form of this command.

map-cache *destination-EID-prefix / prefix-length* {**action** {**drop**| **map-request**| **native-forward**}}| **locator** *locator-address* **priority** *priority-value* **weight** *weight-value*;

no map-cache *destination-EID-prefix / prefix-length* {**action** {**drop**| **map-request**| **native-forward**}}| **locator** *locator-address* **priority** *priority-value* **weight** *weight-value*;

Syntax Description

| | |
|---|--|
| <i>destination-EID-prefix / prefix-length</i> | Specifies the destination IPv4/IPv6 EID-prefix/prefix-length. |
| action | Specifies the non LISP forwarding action for the EID prefix. |
| drop | Selects drop action for the EID-prefix. |
| map-request | Generates a map-request for the EID-prefix. |
| native-forward | Specifies to natively forward EID-prefix. |
| locator <i>locator-address</i> | The IPv4 or IPv6 Routing Locator (RLOC) associated with the EID-prefix/prefix-length. |
| priority <i>priority-value</i> | The priority (value between 0 and 255) assigned to the RLOC. When multiple locators have the same priority they may be used in load-shared fashion. A lower value indicates a higher priority. |
| weight <i>weight-value</i> | The weight (value between 0 and 100) assigned to the locator. Used in order to determine how to load-share traffic between multiple locators when the priorities assigned to multiple locators are the same. The value represents the percentage of traffic to be load-shared. |

Command Default

No IPv6 EID-to-RLOC mapping relationships or static IPv6 EID-to-RLOC mapping destinations are configured.

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

configures a destination EID-to-RLOC mapping and associated traffic policy for the IPv6 EID-prefix block 2001:DB8:BB::/48. The locator for this EID-prefix block is 2001:DB8:0A::1 and the traffic policy for this locator has a priority of 1 and a weight of 100.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#map-cache
RP/0/RSP0/CPU0:router(config-lisp-afi-map-cache)#2001:db8:bb::/48 locator 2001:db8:a::1
priority 1 weight 100
```

Related Commands

| Command | Description |
|--|--|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

map-request-source

To configure inner header source address to be used as the source address for Locator/ID Separation Protocol (LISP) map-request messages, use the **map-request-source** command in LISP address family configuration mode. To remove the configured map-request source address, use the **no** form of this command.

map-request-source *ip_address*

no map-request-source

Syntax Description

| | |
|-------------------|---|
| <i>ip_address</i> | Specifies the IPv4 or IPv6 source address of the inner header in the map-request message. |
|-------------------|---|

Command Default

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure the IPv6 source address to be used as the source address for LISP map-request messages:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#map-request-source 4:5::6
```

Related Commands

| Command | Description |
|---|--|
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

other-xtr-probe

To configure parameters for probing of other xTRs *site-local* routing locators (RLOCs), use the **other-xtr-probe** command in the LISP configuration mode. To return to the default setting, use the **no** form of this command.

other-xtr-probe period *seconds*

no other-xtr-probe period *seconds*

Syntax Description

| | |
|----------------|--|
| period | Configures the site-local RLOC probing period. |
| <i>seconds</i> | Specifies the value of the probing period in seconds. The range is between 5 to 900. |

Command Default

Probing of site-local RLOCs is enabled by default and cannot be disabled. The default interval is 30 seconds.

Command Modes

LISP configuration

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Use the **other-xtr-probe** command to change the probe interval for sending RLOC probes to all site-other RLOCs.

This command probes site-local rlocs, whereas rloc-probing probes remote RLOCs. The latter can be turned on or off with `loc-reach-algorithm rloc-probing`. `Remote-rloc-probe` also probes remote RLOCs.



Note

This functionality is enabled by default and cannot be disabled. The default interval is 30 seconds. Use the **show run|include other-xtr-probe** command to display the configured interval. When an output value is displayed, the value is configured for something other than the default value. When no output is displayed, it is configured for the default.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure the probing period of 86 seconds for probing RLOCs:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#other-xtr-probe period 86
```

Related Commands

| Command | Description |
|--|---|
| remote-rloc-probe , on page 1897 | Configures parameters for probing of remote local routing locators (RLOCs) |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

proxy-etr

To configure a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR), use the **proxy-etr** command in LISP address family configuration mode. To remove LISP PETR functionality, use the **no** form of this command.

proxy-etr

no proxy-etr

This command has no keywords or arguments.

Command Default ETR functionality is disabled.

Command Modes LISP IPv4 address family
LISP IPv6 address family

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.0 | This command was introduced. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | lisp | read, write |

Examples This example shows how to configure PETR functionality under LISP IPv6 on the router:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#proxy-etr
```

| Related Commands | Command | Description |
|------------------|--|---|
| | address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

| Command | Description |
|--|---|
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

proxy-itr

To configure a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR), use the **proxy-itr** command in LISP address family configuration mode. To remove LISP PITR functionality, use the **no** form of this command.

proxy-itr *IPv4-source-locator-address*

noproxy-itr *IPv4-source-locator-address*

Syntax Description

IPv4-source-locator-address

Specifies the IPv4 source locator for PITR.

Command Default

PITR functionality is disabled.

Command Modes

LISP IPv4 address family

LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.0 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure LISP PITR functionality on the router:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-afi)#proxy-itr 10.2.3.4
```

Related Commands

| Command | Description |
|--|---|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

remote-rloc-probe

To configure parameters for probing of remote local routing locators (RLOCs), use the **remote-rloc-probe** command in the LISP configuration mode. To return to the default setting, use the **no** form of this command.

remote-rloc-probe on-route-change

no remote-rloc-probe on-route-change

| Syntax Description | on-route-change | Specifies the probing of the trigger on routing changes for remote RLOCs. |
|--------------------|-----------------|---|
|--------------------|-----------------|---|

| Command Default | None |
|-----------------|------|
|-----------------|------|

| Command Modes | LISP configuration |
|---------------|--------------------|
|---------------|--------------------|

| Command History | Release | Modification |
|-----------------|---------------|------------------------------|
| | Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|---------|---------|-------------|
| | lisp | read, write |

Examples

This example shows how to configure the **remote-rloc-probe** command for remote RLOCs:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#remote-rloc-probe on-route-change
```

| Related Commands | Command | Description |
|------------------|--|---|
| | address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |

| Command | Description |
|---|--|
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

router lisp

To enter Locator and ID Separation Protocol (LISP) configuration mode, use the **router lisp** command in global configuration mode. To remove all LISP configurations and exit the LISP configuration mode, use the **no** form of this command.

router lisp *process-number*

no router lisp

| | | |
|---------------------------|-----------------------|--|
| Syntax Description | <i>process-number</i> | Specifies the router LISP process number. The range is 0-15. |
|---------------------------|-----------------------|--|

Syntax Description This command has no keywords or arguments.

Command Default LISP configuration is disabled.

Command Modes global configuration

| Command History | Release | Modification |
|------------------------|----------------|---|
| | Release 4.3.0 | This command was introduced. |
| | Release 4.3.1 | Support for LISP <i>process-number</i> was added. |

Usage Guidelines To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

| Task ID | Task ID | Operation |
|----------------|----------------|------------------|
| | lisp | read, write |

Examples This example shows how to configure LISP configuration mode:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#
```

Related Commands

| Command | Description |
|--|---|
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| loc-reach-algorithm rloc-probing, on page 1883 | Configures locator reachability algorithm, RLOC Probing, and determines reachability status for other iBGP peers through the IGP domain. |
| map-cache-limit, on page 1885 | Configures the maximum limit of IPv4 LISP or IPv6 LISP map-cache entries allowed to be stored by the router. |
| map-cache, on page 1887 | Configures a static IPv4 EID-to-RLOC or static IPv6 EID-to-RLOC mapping relationship and its associated traffic policy, or statically configures the packet handling behavior associated with a destination IPv4 EID-prefix or a destination IPv6 EID-prefix. |
| proxy-etr, on page 1893 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Egress Tunnel Router (PETR). |
| proxy-itr, on page 1895 | Configures a router to act as an IPv4 or IPv6 LISP Proxy Ingress Tunnel Router (PITR). |

solicit-map-request

To configure the solicit map request (SMR) handling, use the **solicit-map-request** command in the LISP address family configuration mode. To disable solicit map request handling, use the **no** form of this command.

solicit-map-request {**ignore**| **max-per-entry** *number*| **suppression-time** *seconds*}

no solicit-map-request {**ignore**| **max-per-entry** *number*| **suppression-time** *seconds*}

Syntax Description

| | |
|-------------------------|---|
| ignore | Ignores an IPv4 or IPv6 map-request message that has the solicit-map-request (SMR) bit set. |
| max-per-entry | Specifies the maximum number of solicit-map-requests (SMRs) for addresses under a map-cache entry. |
| <i>number</i> | Specifies the maximum number of SMRs. The value ranges from 1 to 100. |
| suppression-time | Specifies how long to suppress repeated solicit-map-requests (SMRs) for the same address. |
| <i>seconds</i> | Specifies the seconds to suppress repeated SMRs for the same address. The value ranges from 1 to 600. |

Command Default

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to configure the **solicit-map-request** command to ignore a map-request message, to specify a maximum number of 50 SMRs, and to specify a suppression time of 500 s:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request ignore
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request max-per-entry 50
RP/0/RSP0/CPU0:router(config-lisp-af)#solicit-map-request suppression-time 500
```

Related Commands

| Command | Description |
|---|--|
| itr map-resolver, on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| etr, on page 1867 | Enables Egress Tunnel Router (ETR) functionality. |
| address-family (LISP), on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| router lisp, on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |

use-petr

To specify that packets hitting a forward-native map-cache entry should be LISP encapsulated and forwarded to a PETR, instead of attempting to forward them natively, use the **use-petr** command in LISP address family configuration mode. To revert to the default behavior of forwarding packets natively, if they hit a forward-native map cache entry, use the **no** form of this command.

use-petr *ip_address* [**priority** *priority* **weight** *weight*]

no use-petr

Syntax Description

| | |
|-------------------|--|
| <i>ip_address</i> | Specifies the IPv4 or IPv6 locator address of the PETR. |
| priority | (Optional) Specifies the priority assigned to this PETR. |
| <i>priority</i> | Specifies the value of the priority assigned to this PETR. This value ranges from 0 to 255. A lower value indicates a higher priority. |
| weight | (Optional) Specifies the percentage of traffic to be load-shared. |
| <i>weight</i> | Specifies the weight in value of the percentage of traffic to be load-shared. The value ranges from 0 to 100. |

Command Default

PETR services are disabled by default.

Command Modes

LISP IPv4 address family
LISP IPv6 address family

Command History

| Release | Modification |
|---------------|------------------------------|
| Release 4.3.1 | This command was introduced. |

Usage Guidelines

To use this command, you must be in a user group associated with a task group that includes appropriate task IDs. If the user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Multiple PETRs can be configured. If you configure a priority and weight for one of them, you must configure a priority and weight for all of the PETRs.

Task ID

| Task ID | Operation |
|---------|-------------|
| lisp | read, write |

Examples

This example shows how to use a single PETR:

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)#use-petr 2001:db8::1
```

This example shows how to configure a fallback PETR. This type of configuration is useful if, by default, you want to reach your PETR over IPv6, but in case you lose IPv6 connectivity to the outside world, you fall back to using IPv4. That way, your IPv6 EID prefix will maintain IPv6 connectivity to the outside world regardless of whether or not you have IPv6 connectivity at the ITR.

```
RP/0/RSP0/CPU0:router#configure
RP/0/RSP0/CPU0:router(config)#router lisp
RP/0/RSP0/CPU0:router(config-lisp)#address-family ipv6 unicast
RP/0/RSP0/CPU0:router(config-lisp-af)# use-petr 2001:db8::1 priority 1 weight 100
RP/0/RSP0/CPU0:router(config-lisp-af)#use-petr 64.10.10.1 priority 2 weight 100
```

Related Commands

| Command | Description |
|--|--|
| itr map-resolver , on page 1875 | Configures an IPv4 or IPv6 locator address of the LISP Map-Resolver to be used by the ITR. |
| address-family (LISP) , on page 1861 | Enters Locator ID and separation protocol (LISP) address family configuration mode. |
| router lisp , on page 1899 | Enters Locator and ID Separation Protocol (LISP) configuration mode. |



INDEX

A

abort (RPL) command [1520](#)
accept-own command [8](#)
add command [1522](#)
additional-paths install backup command [9](#)
additional-paths selection command [13](#)
additional-paths send command [11, 15](#)
address-family (BGP) command [17](#)
address-family (EIGRP) command [605](#)
address-family (IS-IS) command [707](#)
address-family (LISP) command [1861](#)
address-family (OSPF) command [931](#)
address-family (OSPFv3) command [1174](#)
address-family (static) command [1810](#)
address-family ipv4 unicast (BFD) command [527](#)
address-family multicast topology (IS-IS) command [709](#)
address-family next-hop dampening disable command [1355](#)
adjacency stagger command [932](#)
adjacency-check disable command [711](#)
advertise best-external command [21](#)
advertisement-interval command [24](#)
af-group command [26](#)
aggregate-address command [28](#)
aigp command [31](#)
aigp send-cost-community command [33](#)
allocate-label command [35](#)
allowas-in command [37](#)
apply command [1524](#)
area (OSPF) command [934](#)
area (OSPFv3) command [1175](#)
as-format command [39](#)
as-override command [40](#)
as-path in command [1526](#)
as-path is-local command [1528](#)
as-path length command [1530](#)
as-path neighbor-is command [1532](#)
as-path originates-from command [1535](#)
as-path passes-through command [1538](#)
as-path unique-length command [1543](#)
as-path-loopcheck out disable command [42](#)
as-path-set command [1541](#)

authentication (OSPF) command [936](#)
authentication (OSPFv3) command [1177](#)
authentication keychain command [607](#)
authentication keychain mode md5 (RIP) command [1445](#)
authentication-key (OSPF) command [939](#)
auto-cost (OSPF) command [941](#)
auto-cost (OSPFv3) command [1179](#)
auto-summary (EIGRP) command [609](#)
auto-summary (RIP) command [1447](#)
autonomous-system command [611](#)

B

bandwidth-percent (EIGRP) command [613](#)
bfd (BGP) command [46](#)
bfd address-family ipv4 destination command [531](#)
bfd address-family ipv4 fast-detect command [533](#)
bfd address-family ipv4 minimum-interval command [534](#)
bfd address-family ipv4 multiplier command [537](#)
bfd address-family ipv4 timers command [540](#)
bfd command [529](#)
bfd fast-detect command [546](#)
bfd minimum-interval command [549](#)
bfd multiplier command [555](#)
bgp as-path-loopcheck command [51](#)
bgp attribute-download command [52](#)
bgp auto-policy-soft-reset disable command [54](#)
bgp bestpath as-path ignore command [56](#)
bgp bestpath compare-routerid command [58](#)
bgp bestpath cost-community ignore command [60](#)
bgp bestpath med always command [62](#)
bgp bestpath med confed command [64](#)
bgp bestpath med missing-as-worst command [66](#)
bgp client-to-client reflection disable command [68](#)
bgp cluster-id command [70](#)
bgp confederation identifier command [72](#)
bgp confederation peers command [74](#)
bgp dampening command [76](#)
bgp default local-preference command [78](#)
bgp enforce-first-as disable command [79](#)
bgp fast-external-falover disable command [81](#)

bgp graceful-restart command [82](#)
 bgp graceful-restart graceful-reset command [84](#)
 bgp graceful-restart purge-time command [86](#)
 bgp graceful-restart restart-time command [88](#)
 bgp graceful-restart stalepath-time command [90](#)
 bgp import-delay command [92](#)
 bgp label-delay command [94](#)
 bgp log neighbor changes disable command [96](#)
 bgp maximum neighbor command [98](#)
 bgp redistribute-internal command [99](#)
 bgp router-id command [101](#)
 bgp scan-time command [103](#)
 bgp update-delay command [105](#)
 bgp write-limit command [107](#)
 broadcast-for-v2 command [1449](#)
 bundle coexistence bob-blb command [542](#)

C

capability additional-paths receive command [109](#)
 capability additional-paths send command [111](#)
 capability opaque disable command [943](#)
 capability orf prefix command [113](#)
 capability suppress 4-byte-as command [116](#)
 capability vrf-lite (OSPFv3) command [1181](#)
 circuit-type command [713](#)
 clear bfd counters command [558](#)
 clear bgp command [119](#)
 clear bgp dampening command [121](#)
 clear bgp external command [123](#)
 clear bgp flap-statistics command [125](#)
 clear bgp nexthop performance-statistics command [127](#)
 clear bgp nexthop registration command [129](#)
 clear bgp peer-drops command [131](#)
 clear bgp performance-statistics command [132](#)
 clear bgp self-originated command [133](#)
 clear bgp shutdown command [135](#)
 clear bgp soft command [137](#)
 clear eigrp neighbors command [615](#)
 clear eigrp topology command [617](#)
 clear isis process command [715](#)
 clear isis route command [717](#)
 clear isis statistics command [719](#)
 clear ospf process command [946](#)
 clear ospf redistribution command [948](#)
 clear ospf routes command [950](#)
 clear ospf statistics command [952](#)
 clear ospfv3 process command [1183](#)
 clear ospfv3 redistribution command [1185](#)
 clear ospfv3 routes command [1187](#)
 clear ospfv3 statistics command [1189](#)
 clear rip command [1450](#)

clear rip database command [1452](#)
 clear rip interface command [1454](#)
 clear rip out-of-memory command [1456](#)
 clear rip statistics command [1458](#)
 clear route command [1356](#)
 collect-diagnostics (RCMD) command [1829](#)
 community is-empty command [1545](#)
 community matches-any command [1546](#)
 community matches-every command [1548](#)
 community-set command [1551](#)
 cost (OSPF) command [954](#)
 cost (OSPFv3) command [1191](#)
 cost-fallback (OSPF) command [956](#)
 csnp-interval command [721](#)

D

dampening (BFD Session) command [544](#)
 dampening (BFD) command [560](#)
 database-filter all out (OSPF) command [958](#)
 database-filter all out (OSPFv3) command [1193](#)
 dead-interval (OSPF) command [960](#)
 dead-interval (OSPFv3) command [1195](#)
 default-cost (OSPF) command [962](#)
 default-cost (OSPFv3) command [1197](#)
 default-information command [619](#)
 default-information originate (BGP) command [140](#)
 default-information originate (IS-IS) command [723](#)
 default-information originate (OSPF) command [964](#)
 default-information originate (OSPFv3) command [1199](#)
 default-information originate (RIP) command [1460](#)
 default-metric (BGP) command [143](#)
 default-metric (EIGRP) command [621](#)
 default-metric (OSPF) command [966](#)
 default-metric (OSPFv3) command [1201](#)
 default-metric (RIP) command [1462](#)
 default-originate command [144](#)
 delete community command [1554](#)
 delete extcommunity rt command [1556](#)
 demand-circuit (OSPF) command [968](#)
 demand-circuit (OSPFv3) command [1203](#)
 description (BGP) command [146](#)
 destination in command [1558](#)
 disable (IS-IS) command [725](#)
 disable-dn-bit-check command [970](#)
 distance (EIGRP) command [623](#)
 distance (IS-IS) command [726](#)
 distance (OSPF) command [971](#)
 distance (RIP) command [1464](#)
 distance bgp command [148](#)
 distance ospf command [974](#)
 distance ospfv3 command [1205](#)

distribute-list command [976](#)
 distribute-list prefix-list in command [1207](#)
 distribute-list prefix-list out command [1209](#)
 dmz-link-bandwidth command [154](#)
 domain-id (OSPF) command [978](#)
 domain-id (OSPFv3) command [1212](#)
 domain-tag command [980](#)
 done command [1560](#)
 drop command [1562](#)
 dscp (BGP) command [156](#)

E

ebgp-multihop command [159](#)
 echo disable command [563](#)
 echo ipv4 source command [565](#)
 echo latency detect command [567](#)
 echo startup validate command [569](#)
 edit command [1564](#)
 eid-mtu command [1863](#)
 eid-table command [1865](#)
 encryption command [1214](#)
 end-global command [1567](#)
 end-policy command [1568](#)
 end-set command [1570](#)
 etr accept-map-request-mapping command [1869](#)
 etr command [1867](#)
 etr map-cache-ttl command [1871](#)
 etr map-server command [1873](#)
 event-buffer-size (RCMD) command [1831](#)
 export route-policy command [161](#)
 export route-target command [163](#)
 extcommunity rt is-empty command [1572](#)
 extcommunity rt matches-any command [1574](#)
 extcommunity rt matches-every command [1576](#)
 extcommunity rt matches-within command [1578](#)
 extcommunity soo is-empty command [1586](#)
 extcommunity soo matches-any command [1588](#)
 extcommunity soo matches-every command [1590](#)
 extcommunity-set cost command [1580](#)
 extcommunity-set rt command [1582](#)
 extcommunity-set soo command [1584](#)

F

fast-reroute (OSPFv2) command [981](#)
 fast-reroute per-link (IS-IS) command [728](#)
 fast-reroute per-link priority-limit command [732](#)
 fast-reroute per-prefix (IS-IS) command [730](#)
 fast-reroute per-prefix load-sharing disable command [734](#)
 fast-reroute per-prefix tiebreaker command [735](#)

flood-reduction (OSPF) command [991](#)
 flood-reduction (OSPFv3) command [1216](#)

G

graceful-restart (OSPFv3) command [1218](#)

H

hello-interval (EIGRP) command [625](#)
 hello-interval (IS-IS) command [737](#)
 hello-interval (OSPF) command [993](#)
 hello-interval (OSPFv3) command [1220](#)
 hello-multiplier command [739](#)
 hello-padding command [741](#)
 hello-password accept command [747](#)
 hello-password command [743](#)
 hello-password keychain command [745](#)
 hold-time (EIGRP) command [626](#)
 hostname dynamic disable command [749](#)

I

if command [1592](#)
 ignore lsa mospf command [995](#)
 ignore-connected-check command [169](#)
 ignore-lsp-errors command [751](#)
 import route-policy command [165](#)
 import route-target command [167](#)
 instance command [1222](#)
 interface (BFD) command [571](#)
 interface (EIGRP) command [628](#)
 interface (IS-IS) command [753](#)
 interface (OSPF) command [996](#)
 interface (OSPFv3) command [1224](#)
 interface (RIP) command [1467](#)
 ipfrr lfa exclude interface command [983](#)
 ipv6 checksum command [574](#)
 is-type command [756](#)
 ispf command [755](#)
 itr map-resolver command [1875](#)

K

keychain command [170](#)
 keychain inheritance-disable command [174](#)
 keychain-disable command [172](#)

L

label mode command [178](#)
 label-allocation-mode command [176](#)
 lcc command [1360](#)
 loc-reach-algorithm rloc-probing command [1883](#)
 local-as command [181](#)
 local-preference command [1599](#)
 locator reachability command [1877](#)
 locator-set command [1879](#)
 locator-table command [1881](#)
 Locator/ID Separation Protocol (LISP) [1860](#)
 log adjacency changes (IS-IS) command [758](#)
 log adjacency changes (OSPF) command [998](#)
 log adjacency changes (OSPFv3) command [1226](#)
 log pdu drops command [760](#)
 log-neighbor-changes command [630](#)
 log-neighbor-warnings command [631](#)
 loopback stub-network command [999](#)
 lsp fast-flood threshold command [761](#)
 lsp-gen-interval command [763](#)
 lsp-interval command [765](#)
 lsp-mtu command [767](#)
 lsp-password accept command [771](#)
 lsp-password command [769](#)
 lsp-refresh-interval command [773](#)

M

map-cache command [1887](#)
 map-cache-limit command [1885](#)
 map-request-source command [1889](#)
 max-events-stored (RCMD) command [1833](#)
 max-lsa command [1001](#)
 max-lsp-lifetime command [778](#)
 max-metric command [1004](#)
 maximum interfaces (OSPF) command [1007](#)
 maximum interfaces (OSPFv3) command [1228](#)
 maximum path (static) command [1812](#)
 maximum paths (OSPF) command [1009](#)
 maximum paths (OSPFv3) command [1230](#)
 maximum prefix (RIB) command [1358](#)
 maximum redistributed-prefixes (OSPF) command [1011](#)
 maximum redistributed-prefixes (OSPFv3) command [1232](#)
 maximum-paths (BGP) command [183](#)
 maximum-paths (EIGRP) command [632](#)
 maximum-paths (IS-IS) command [775](#)
 maximum-paths (RIP) command [1469](#)
 maximum-prefix (BGP) command [185](#)
 maximum-prefix (EIGRP) command [634](#)
 maximum-redistributed-prefixes (IS-IS) command [776](#)
 med command [1601](#)
 mesh-group (IS-IS) command [780](#)

message-digest-key command [1013](#)
 metric (EIGRP) command [637](#)
 metric command [782](#)
 metric maximum-hops command [639](#)
 metric weights command [643](#)
 metric-style narrow command [784](#)
 metric-style transition command [786](#)
 metric-style wide command [788](#)
 metric-zero-accept command [1470](#)
 min-lsp-arrivaltime command [790](#)
 monitor-convergence (IS-IS) command [1826](#)
 monitor-convergence (OSPF) command [1828](#)
 monitoring-interval (RCMD) command [1835](#)
 mpls activate (BGP) command [188](#)
 mpls ldp auto-config (OSPF) command [1016](#)
 mpls ldp auto-config command [792](#)
 mpls ldp sync (IS-IS) command [794](#)
 mpls ldp sync (OSPF) command [1017](#)
 mpls traffic-eng (OSPF) command [1018](#)
 mpls traffic-eng command [796](#)
 mpls traffic-eng igp-intact (OSPF) command [1020](#)
 mpls traffic-eng multicast-intact (IS-IS) command [798](#)
 mpls traffic-eng multicast-intact (OSPF) command [1022](#)
 mpls traffic-eng path-selection ignore overload command [800](#)
 mpls traffic-eng router-id (IS-IS) command [802](#)
 mpls traffic-eng router-id (OSPF) command [1024](#)
 mtu-ignore (OSPF) command [1026](#)
 mtu-ignore (OSPFv3) command [1234](#)
 multi-area-interface command [1028](#)
 mvpn command [191](#)

N

neighbor (BGP) command [192](#)
 neighbor (EIGRP) command [646](#)
 neighbor (OSPF) command [1030](#)
 neighbor (OSPFv3) command [1236](#)
 neighbor (RIP) command [1471](#)
 neighbor database-filter all out command [1033](#)
 neighbor maximum-prefix command [648](#)
 neighbor-group command [194](#)
 net command [804](#)
 network (BGP) command [196](#)
 network (OSPF) command [1035](#)
 network (OSPFv3) command [1238](#)
 network backdoor command [198](#)
 next-hop in command [1603](#)
 next-hop-self command [200](#)
 next-hop-self disable command [651](#)
 next-hop-unchanged command [203](#)
 nexthop resolution prefix-length minimum command [205](#)
 nexthop route-policy command [206](#)

nexthop trigger-delay command [208](#)
 node disable (RCMD) command [1837](#)
 nsf (IS-IS) command [806](#)
 nsf (OSPF) command [1037](#)
 nsf (RIP) command [1473](#)
 nsf disable command [652](#)
 nsf flush-delay-time (OSPF) command [1039](#)
 nsf interface-expires command [808](#)
 nsf interface-timer command [810](#)
 nsf interval (OSPF) command [1040](#)
 nsf lifetime (IS-IS) command [812](#)
 nsf lifetime (OSPF) command [1042](#)
 nsr (OSPF) command [1044](#)
 nsr (OSPFv3) command [1242](#)
 nsr command [210](#)
 nssa (OSPF) command [1045](#)
 nssa (OSPFv3) command [1240](#)

O

orf command [212](#)
 orf prefix in command [1605](#)
 origin is command [1607](#)
 ospf name-lookup command [1047](#)
 ospfv3 name-lookup command [1244](#)
 other-xtr-probe command [1891](#)
 output-delay command [1474](#)

P

packet-size command [1048, 1245](#)
 pass command [1609](#)
 passive (IS-IS) command [813](#)
 passive (OSPF) command [1050](#)
 passive (OSPFv3) command [1247](#)
 passive-interface (EIGRP) command [654](#)
 passive-interface (RIP) command [1475](#)
 password (BGP) command [214](#)
 password (rpki-cache) command [216](#)
 password-disable command [218](#)
 path-type is command [1611](#)
 point-to-point command [814](#)
 poison-reverse command [1477](#)
 policy-global command [1613](#)
 precedence command [222](#)
 preference (rpki-server) command [224](#)
 prefix-set command [1615](#)
 prepend as-path command [1618](#)
 priority (IS-IS) command [815](#)
 priority (OSPF) command [1052](#)
 priority (OSPFv3) command [1249](#)

priority (RCMD) command [1843](#)
 propagate level command [817](#)
 protocol (RCMD) command [1845](#)
 protocol command [1620](#)
 protocol shutdown command [1054, 1251](#)
 proxy-etr command [1893](#)
 proxy-itr command [1895](#)
 purge-time (rpki-cache) command [225](#)

Q

queue dispatch flush-lsa command [1055](#)
 queue dispatch incoming command [1057](#)
 queue dispatch rate-limited-lsa command [1059](#)
 queue dispatch spf-lsa-limit command [1061](#)
 queue limit command [1063](#)

R

range (OSPF) command [1065](#)
 range (OSPFv3) command [1252](#)
 rcc command [1362](#)
 rd command [227](#)
 rd in command [1622](#)
 rd-set command [1624](#)
 receive version command [1479](#)
 receive-buffer-size command [229](#)
 recursion-depth-max command [1364](#)
 redistribute (BGP) command [231](#)
 redistribute (EIGRP) command [655](#)
 redistribute (IS-IS) command [819](#)
 redistribute (OSPF) command [1067](#)
 redistribute (OSPFv3) command [1254](#)
 redistribute (RIP) command [1481](#)
 redistribute maximum-prefix command [658](#)
 refresh-time (rpki-cache) command [235](#)
 remote-as (BGP) command [239](#)
 remote-rloc-probe [1897](#)
 remove-private-as command [242](#)
 replace as-path command [1626](#)
 response-time (rpki-cache) command [237](#)
 retain local-label command [245](#)
 retain route-target command [247](#)
 retransmit-interval (IS-IS) command [823](#)
 retransmit-interval (OSPF) command [1072](#)
 retransmit-interval (OSPFv3) command [1259](#)
 retransmit-throttle-interval command [825](#)
 rib-has-route command [1628](#)
 route (static) command [1814](#)
 route-has-label command [1630](#)
 route-policy (BGP) command [249](#)

route-policy (EIGRP) command [660](#)
 route-policy (OSPF) command [1074](#)
 route-policy (RIP) command [1487](#)
 route-policy (RPL) command [1631](#)
 route-reflector-client command [251](#)
 route-type is command [1633](#)
 router bgp command [253](#)
 router eigrp command [662](#)
 router isis command [827](#)
 router lisp command [1899](#)
 router ospf command [1078](#)
 router ospfv3 command [1263](#)
 router rib command [1365](#)
 router rip command [1485](#)
 router static command [1818](#)
 router-convergence command [1824](#)
 router-id (EIGRP) command [664](#)
 router-id (OSPF) command [1076](#)
 router-id (OSPFv3) command [1261](#)
 rpki cache command [255](#)
 rpl editor command [1635](#)
 rpl maximum command [1637](#)
 rump always-replicate command [1366](#)

S

security ttl (OSPF) command [1080](#)
 selective-vrf-download disable command [256](#)
 send version command [1489](#)
 send-buffer-size command [258](#)
 send-community-ebgp command [260](#)
 send-extended-community-ebgp command [262](#)
 session-group command [264](#)
 session-open-mode command [266](#)
 set aigp-metric command [1639](#)
 set community command [1641](#)
 set core-tree command [1643](#)
 set dampening command [1645](#)
 set eigrp-metric command [1647](#)
 set extcommunity cost command [1649](#)
 set extcommunity rt command [1651](#)
 set ip-precedence command [1653](#)
 set isis-metric command [1655](#)
 set label command [1656](#)
 set label-mode command [1657](#)
 set level command [1659](#)
 set local-preference command [1661](#)
 set med command [1663](#)
 set metric-type (IS-IS) command [1665](#)
 set metric-type (OSPF) command [1667](#)
 set next-hop command [1669](#)
 set origin command [1671](#)
 set ospf-metric command [1673](#)
 set path-selection command [1674](#)
 set qos-group command [1676](#)
 set rib-metric command [1678](#)
 set rip-metric command [1679](#)
 set rip-tag command [1681](#)
 set rpf-topology command [1682](#)
 set spf-priority command [1684](#)
 set tag command [1686](#)
 set traffic-index command [1688](#)
 set vpn-distinguisher command [1690](#)
 set weight command [1692](#)
 set-attached-bit command [829](#)
 set-overload-bit command [831](#)
 sham-link command [1082, 1265](#)
 show bfd client command [581](#)
 show bfd command [578](#)
 show bfd counters command [583](#)
 show bfd mib session command [586](#)
 show bfd neighbor command [591](#)
 show bfd session command [593](#)
 show bgp advertised command [293](#)
 show bgp af-group command [301](#)
 show bgp attribute-key command [304](#)
 show bgp cidr-only command [309](#)
 show bgp command [268](#)
 show bgp community command [313](#)
 show bgp convergence command [318](#)
 show bgp dampened-paths command [321](#)
 show bgp flap-statistics command [326](#)
 show bgp inconsistent-as command [332](#)
 show bgp l2vpn rd command [341](#)
 show bgp labels command [337](#)
 show bgp neighbor-group command [349](#)
 show bgp neighbors command [353](#)
 show bgp neighbors nsr command [375](#)
 show bgp nexthops command [377](#)
 show bgp nsr command [386](#)
 show bgp paths command [390](#)
 show bgp policy command [393](#)
 show bgp process command [402](#)
 show bgp regexp command [426](#)
 show bgp route-policy command [431](#)
 show bgp session-group command [436](#)
 show bgp sessions command [439](#)
 show bgp summary command [442](#)
 show bgp summary nsr command [447](#)
 show bgp truncated-communities command [451](#)
 show bgp update out command [282](#)
 show bgp update out process command [286](#)
 show bgp update-group command [456](#)
 show bgp vrf imported-routes command [464](#)
 show eigrp accounting command [665](#)
 show eigrp interfaces command [668](#)

- show eigrp neighbors command [673](#)
- show eigrp topology command [677](#)
- show eigrp traffic command [681](#)
- show isis adjacency command [836](#)
- show isis adjacency-log command [839](#)
- show isis checkpoint adjacency command [841](#)
- show isis checkpoint interface command [843](#)
- show isis checkpoint lsp command [845](#)
- show isis command [833](#)
- show isis database command [847](#)
- show isis database-log command [856](#)
- show isis fast-reroute command [859](#)
- show isis hostname command [861](#)
- show isis interface command [863](#)
- show isis lsp-log command [868](#)
- show isis mesh-group command [871](#)
- show isis mpls traffic-eng adjacency-log command [873](#)
- show isis mpls traffic-eng advertisements command [875](#)
- show isis mpls traffic-eng tunnel command [878](#)
- show isis neighbors command [880](#)
- show isis protocol command [884](#)
- show isis route command [887](#)
- show isis spf-log command [890](#)
- show isis statistics command [898](#)
- show isis topology command [902](#)
- show lcc statistics command [1368](#)
- show ospf border-routers command [1088](#)
- show ospf command [1084](#)
- show ospf database command [1090](#)
- show ospf flood-list command [1104](#)
- show ospf interface command [1106](#)
- show ospf message-queue command [1114](#)
- show ospf mpls traffic-eng command [1109](#)
- show ospf neighbor command [1117](#)
- show ospf request-list command [1125](#)
- show ospf retransmission-list command [1128](#)
- show ospf routes command [1131](#)
- show ospf sham-links command [1136](#)
- show ospf summary-prefix command [1139](#)
- show ospf virtual-links command [1141](#)
- show ospfv3 border-routers command [1272](#)
- show ospfv3 command [1267](#)
- show ospfv3 database command [1274](#)
- show ospfv3 flood-list command [1287](#)
- show ospfv3 interface command [1290](#)
- show ospfv3 message-queue command [1293](#)
- show ospfv3 neighbor command [1295](#)
- show ospfv3 request-list command [1302](#)
- show ospfv3 retransmission-list command [1305](#)
- show ospfv3 routes command [1308](#)
- show ospfv3 statistics rib-thread command [1311](#)
- show ospfv3 summary-prefix command [1313](#)
- show ospfv3 virtual-links command [1315](#)
- show protocols (BGP) command [467](#)
- show protocols (EIGRP) command [683](#)
- show protocols (IS-IS) command [905](#)
- show protocols (OSPF) command [1143](#)
- show protocols (OSPFv3) command [1317](#)
- show protocols (RIP) command [1491](#)
- show rcc command [1370](#)
- show rib afi-all command [1378](#)
- show rib attributes command [1380](#)
- show rib client-id command [1382](#)
- show rib clients command [1384](#)
- show rib command [1376](#)
- show rib extcomms command [1386](#)
- show rib firsthop command [1388](#)
- show rib history command [1390](#)
- show rib next-hop command [1392](#)
- show rib opaques command [1394](#)
- show rib protocols command [1397](#)
- show rib recursion-depth-max command [1399](#)
- show rib statistics command [1401](#)
- show rib tables command [1404](#)
- show rib trace command [1406](#)
- show rib vpn-attributes command [1408](#)
- show rib vrf command [1410](#)
- show rip command [1493](#)
- show rip database command [1495](#)
- show rip interface command [1497](#)
- show rip statistics command [1504](#)
- show route backup command [1418](#)
- show route best-local command [1421](#)
- show route command [1412](#)
- show route connected command [1423](#)
- show route local command [1425](#)
- show route longer-prefixes command [1427](#)
- show route next-hop command [1430](#)
- show route quarantined command [1433](#)
- show route resolving-next-hop command [1436](#)
- show route static command [1438](#)
- show route summary command [1440](#)
- show rpl active as-path-set command [1696](#)
- show rpl active community-set command [1699](#)
- show rpl active extcommunity-set command [1702](#)
- show rpl active prefix-set command [1705](#)
- show rpl active rd-set command [1708](#)
- show rpl active route-policy command [1710](#)
- show rpl as-path-set attachpoints command [1715](#)
- show rpl as-path-set command [1713](#)
- show rpl as-path-set references command [1718](#)
- show rpl command [1694](#)
- show rpl community-set attachpoints command [1723](#)
- show rpl community-set command [1721](#)
- show rpl community-set references command [1726](#)
- show rpl extcommunity-set command [1729](#)
- show rpl inactive as-path-set command [1732](#)
- show rpl inactive community-set command [1735](#)

show rpl inactive extcommunity-set command [1738](#)
 show rpl inactive prefix-set command [1741](#)
 show rpl inactive rd-set command [1744](#)
 show rpl inactive route-policy command [1746](#)
 show rpl maximum command [1749](#)
 show rpl policy-global references command [1751](#)
 show rpl prefix-set attachpoints command [1755](#)
 show rpl prefix-set command [1753](#)
 show rpl prefix-set references command [1758](#)
 show rpl rd-set attachpoints command [1763](#)
 show rpl rd-set command [1761](#)
 show rpl rd-set references command [1765](#)
 show rpl route-policy attachpoints command [1771](#)
 show rpl route-policy command [1768](#)
 show rpl route-policy inline command [1774](#)
 show rpl route-policy references command [1776](#)
 show rpl route-policy uses command [1779](#)
 show rpl unused as-path-set command [1782](#)
 show rpl unused community-set command [1785](#)
 show rpl unused extcommunity-set command [1788](#)
 show rpl unused prefix-set command [1790](#)
 show rpl unused rd-set command [1793](#)
 show rpl unused route-policy command [1795](#)
 show svd role command [470](#)
 show svd state command [472](#)
 shutdown (BGP) command [473](#)
 shutdown (IS-IS) command [909](#)
 shutdown (rpki-server) command [475](#)
 single-topology command [910](#)
 site-of-origin (BGP) command [477](#)
 site-of-origin (EIGRP) command [687](#)
 site-of-origin (RIP) command [1506](#)
 snmp context (OSPF) command [1146](#)
 snmp context (OSPFv3) command [1319](#)
 snmp trap (OSPF) command [1149](#)
 snmp trap rate-limit (OSPF) command [1150](#)
 snmp trap rate-limit (OSPFv3) command [1323](#)
 snmp-server traps isis command [912](#)
 socket receive-buffer-size command [479](#)
 socket send-buffer-size command [481](#)
 soft-reconfiguration inbound command [483](#)
 solicit-map-request command [1901](#)
 source in command [1798](#)
 speaker-id command [486](#)
 spf prefix-priority (IS-IS) command [915](#)
 spf prefix-priority (OSPFv2) command [1151](#)
 spf prefix-priority (OSPFv3) command [1325](#)
 spf-interval command [913](#)
 split-horizon disable (EIGRP) command [689](#)
 split-horizon disable (RIP) command [1508](#)
 storage-location (RCMD) command [1854](#)
 stub (EIGRP) command [690](#)
 stub (OSPF) command [1153](#)
 stub (OSPFv3) command [1327](#)

stub-router command [1329](#)
 summary-address (EIGRP) command [693](#)
 summary-prefix (IS-IS) command [917](#)
 summary-prefix (OSPF) command [1155](#)
 summary-prefix (OSPFv3) command [1331](#)
 suppress-route command [1800](#)
 suppressed command [919](#)

T

table-policy command [489](#)
 tag (IS-IS) command [921](#)
 tag command [1802](#)
 timers (BGP) command [491](#)
 timers active-time command [695](#)
 timers basic command [1510](#)
 timers bgp command [493](#)
 timers lsa arrival command [1333](#)
 timers lsa group-pacing command [1157](#)
 timers lsa min-arrival command [1159](#)
 timers nsf route-hold (EIGRP) command [697](#)
 timers pacing flood command [1335](#)
 timers pacing lsa-group command [1337](#)
 timers pacing retransmission command [1339](#)
 timers throttle lsa all (OSPF) command [1160](#)
 timers throttle lsa all (OSPFv3) command [1341](#)
 timers throttle spf (OSPF) command [1163](#)
 timers throttle spf (OSPFv3) command [1343](#)
 topology-id command [923](#)
 trace (IS-IS) command [925](#)
 trace (OSPFv3) command [1345](#)
 transmit-delay (OSPF) command [1165](#)
 transmit-delay (OSPFv3) command [1347](#)
 transport (rpki-server) command [495](#)
 ttl-security command [497](#)

U

unsuppress-route command [1804](#)
 update in error-handling basic disable command [506](#)
 update in error-handling extended command [508](#)
 update limit address-family command [502](#)
 update limit command [500](#)
 update limit sub-group command [504](#)
 update out logging command [510](#)
 update wait-install command [513](#)
 update-source command [511](#)
 use command [514](#)
 use-petr command [1903](#)
 username (rpki-server) command [519](#)

V

validate-update-source disable command [1512](#)
variance command [699](#)
virtual-link (OSPF) command [1167](#)
virtual-link (OSPFv3) command [1349](#)
vpn-distinguisher is command [1806](#)
vrf (BGP) command [520](#)
vrf (EIGRP) command [701](#)
vrf (OSPF) command [1169](#)

vrf (OSPFv3) command [1351](#)
vrf (RIP) command [1513](#)
vrf (static) command [1820](#)

W

weight command [521](#)

