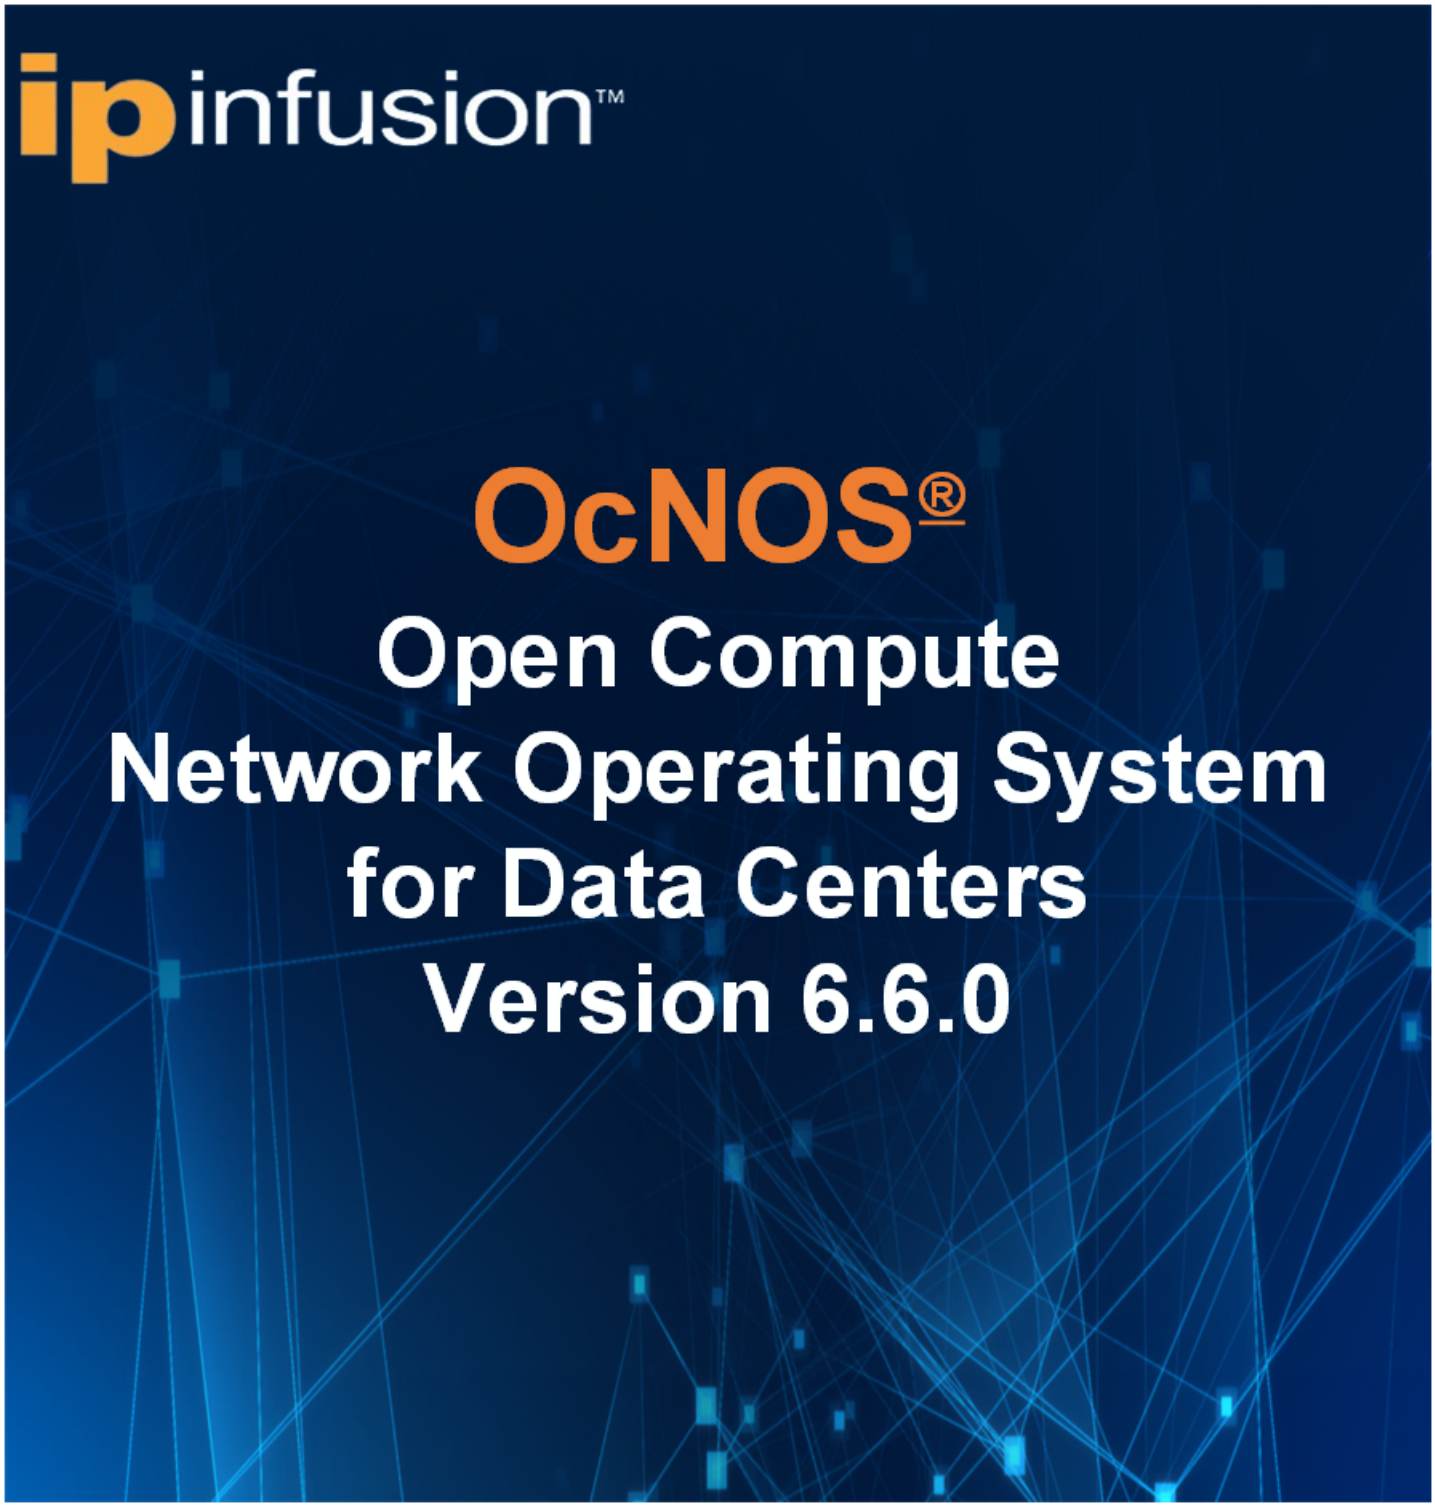

Layer 3 Guide

February 2025

IP Infusion Inc. Proprietary

© 2025 IP Infusion Inc. All Rights Reserved.

This documentation is subject to change without notice. The software described in this document and this documentation are furnished under a license agreement or nondisclosure agreement. The software and documentation may be used or copied only in accordance with the terms of the applicable agreement. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's internal use without the written permission of IP Infusion Inc.

IP Infusion Inc.

3979 Freedom Circle

Suite 900

Santa Clara, California 95054

http://www.ipinfusion.com/

For support, questions, or comments via E-mail, contact:

support@ipinfusion.com

Trademarks:

IP Infusion and OcNOS are trademarks or registered trademarks of IP Infusion. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Use of certain software included in this equipment is subject to the IP Infusion, Inc. End User License Agreement at http://www.ipinfusion.com/license. By using the equipment, you accept the terms of the End User License Agreement.

Contents

Preface 26
IP Maestro Support 26
Audience 26
Conventions 26
Chapter Organization 26
Related Documentation 26
Feature Availability 27
Migration Guide 27
Support 27
Comments 27

Command Line Interface 28
Overview 28
Command Line Interface Help 28
Command Completion 29
Command Abbreviations 29
Command Line Errors 29
Command Negation 30
Syntax Conventions 30
Variable Placeholders 31
Command Description Format 32
Keyboard Operations 32
Show Command Modifiers 33
String Parameters 36
Command Modes 36
Transaction-based Command-line Interface 38

Unicast Configuration 39

CHAPTER 1 Static Routes 40
Topology 40
Configuration 40

CHAPTER 2 Static Route Discard Configuration 45
Overview 45
IPv4 Route Discard 45
IPv6 Route Discard 49

CHAPTER 3 Static Route Object Tracking using IP SLA 54
Overview 54
Topology 54
IPv4 Configuration 55
IPv6 Configuration 61

CHAPTER 4 Route Monitor 65
Overview 65

© 2025 IP Infusion Inc. Proprietary

Prerequisites 65
Configuration 66
Implementation Examples 73
New CLI Commands 73
Troubleshooting 74
Abbreviations 74
Glossory 74

CHAPTER 5 Route-map Continue Configuration 76
Overview 76
Topology 76
Validation 79

CHAPTER 6 Unidirectional Link Detection Configuration 83
Overview 83

CHAPTER 7 Internet Protocol SLA Configuration 86
Configuration 86
Validation 87

CHAPTER 8 Traffic Mirroring using ERSPAN 89
Overview 89
Prerequisites 89
Configuration 90
Validation 92
CLI Commands 92
Glossary 100

Fundamental Layer 3 Command Reference 101

CHAPTER 1 Fundamental Layer 3 Commands 102
automatic-router-id-selection enable 103
clear ip route kernel 104
clear ip route 105
clear ip route vrf NAME 106
clear router-id 107
debug rib 108
description 110
fib retain 111
ip route 112
ip urpf enable 115
ip verify unicast source reachable-via 116
ip vrf 117
ipv6 route 118
maximum-paths 121
max-static-routes 122
router-id 123
show debugging rib 124
show ip rpf 125
show ipv6 rpf 127

© 2025 IP Infusion Inc. Proprietary

show router-id 129
show running-config router 130
show running-config router-id 131
show running-config urpf 132
show running-config vrf 133
snmp restart rib 134

CHAPTER 2 Route-Map Commands 135
continue 137
match as-path 139
match community 140
match extcommunity 141
match interface 142
match ip address 143
match ip address prefix-list 144
match ip next-hop 145
match ip next-hop prefix-list 146
match ip peer 147
match ipv6 address 148
match ipv6 address prefix-list 149
match ipv6 next-hop 150
match ipv6 next-hop prefix-list 151
match ipv6 peer 152
match metric 153
match origin 154
match route-type 155
match tag 156
route-map 157
set aggregator 158
set as-path 159
set atomic-aggregate 160
set comm-list 161
set community 162
set dampening 164
set extcommunity 165
set interface null0 167
set ip next-hop 168
set ipv6 next-hop 169
set level 170
set local-preference 171
set metric 172
set metric-type 173
set origin 174
set originator-id 175
set tag 176
set weight 177
show route-map 178

© 2025 IP Infusion Inc. Proprietary

show running-config route-map 179

CHAPTER 3 Unidirectional Link Detection Commands 180
udld 181
udld message-time 182
udld mode 183
udld state 184
show udld 185
show udld interface 186

CHAPTER 4 Object Tracking Commands 187
track ip sla reachability 188
delay up down 189
show track 190
show track <1-500> 191
show track summary 192
show running-config track 193

CHAPTER 5 IP Service Level Agreements Commands 194
clear ip sla statistics 195
frequency 196
icmp-echo 197
ip sla 198
ip sla schedule 199
show ip sla statistics 200
show ip sla summary 202
show running-config ip sla 203
threshold 204
timeout 205

CHAPTER 6 Source Interface Commands 206
ip source-interface 207
ipv6 source-interface 208
show ip source-interface detail 209
show ipv6 source-interface detail 210
show running-config ip source-interface 211
show running-config ipv6 source-interface 212

Border Gateway Protocol Configuration 213

CHAPTER 1 BGP 214
Enable BGP Routers in the Same Autonomous System 214
Enable BGP Between Different Autonomous Systems 215
Route-Map 217
Route Reflector 219
Multiple Route Reflectors 223
BGP Confederations 234
Dynamic BGP Peering 247
Enable eBGP Multihop 271
Enable Peer Groups 275

Route Redistribution in BGP 286
Add Multiple Instances of the Same Autonomous System 287
Remove the Multi-Exit Disc Attribute from Update Messages 290
Removing Sent and Received MED values 291
BGP Four-Byte Autonomous System 292
4-Octet ASN Capability Enabled on R1 and R2 293
4-Octet ASN Capability Enabled on R1 and Disabled on R2 294
Nexthop Tracking 304
Nexthop Tracking Delay Timer 307
BGP Distance 313
BGP Weight per Peer 317
OSPF as PE-CE Protocol for VPNs 320
BGP Multipath for IPv4 323
Multipath eBGP 326
Multipath eiBGP 329
6BGP AS-PATH Multipath-relax 334
BGP FIB Install (Selective Route Download) 337
BGP Best Path Selection Process 343
BGP Dampening 345
BGP Authentication 351
BGP Unnumbered 353
BGP Blackhole Community Attribute 362

CHAPTER 2 BGP Graceful Restart Configuration 367

CHAPTER 3 BGP Additional Path 375
Overview 375
Prerequisites 375
Configuration 376
Additional Paths at the Global Level 380
Additional Paths Send and Receive at Address-family level 386
Additional Paths at the Neighbor Level 386
Selection of all Additional Paths at the Address-family Level 393
Selection of all Additional Paths at the Neighbor Level 394
Selection of Best 2 Additional Paths at AF Level 396
Selection of Best 2 Additional Paths at the Neighbor Level 396
Selection of Best 3 Additional Paths at the AF Level 398
Selection of Best 3 Additional Paths at the Neighbor Level 399
Implementation Examples 401
CLI Commands 401
Troubleshooting 405
Abbreviations 405

CHAPTER 4 Hide the Remote AS using the neighbor local-as Command 407
Overview 407
Configuration 407
Revised CLI Commands 411
Abbreviations 411

© 2025 IP Infusion Inc. Proprietary

CHAPTER 5 TCP MSS configuration for BGP neighbors 412
Overview 412
Prerequisites 413
Configuration 413
New CLI Commands 417
Abbreviations 418
Glossary 418

CHAPTER 6 BGP Bogon Prefix Filtering IPv4 420
Overview 420
Prerequisites 420
Configuration 422
BGP Bogon Prefix Filtering Command 423
Glossary 424

Border Gateway Protocol Command Reference 425

CHAPTER 1 BGP Commands 426
address-family 430
aggregate-address 432
auto-summary 433
bgp aggregate-nexthop-check 434
bgp always-compare-med 435
bgp as-local-count 436
bgp bestpath as-path ignore 437
bgp bestpath as-path multipath-relax 438
bgp bestpath compare-confed-aspath 439
bgp bestpath compare-routerid 440
bgp bestpath dont-compare-originator-id 441
bgp bestpath med 442
bgp bestpath tie-break-on-age 444
bgp client-to-client reflection 445
bgp cluster-id 446
bgp confederation vrf identifier 447
bgp confederation vrf peers 448
bgp dampening 449
bgp default local-preference 451
bgp deterministic-med 452
bgp enforce-first-as 453
bgp extended-asn-cap 454
bgp fast-external-failover 455
bgp graceful-restart 456
bgp g-shut 458
bgp g-shut-capable 459
bgp g-shut-local-preference 460
bgp log-neighbor-changes 461
bgp nexthop-trigger delay 463
bgp nexthop-trigger enable 464

© 2025 IP Infusion Inc. Proprietary

bgp rfc1771-path-select 465
bgp rfc1771-strict 466
bgp router-id 467
bgp scan-time 468
bgp table-map 469
bgp update-delay 471
clear bgp (A.B.C.D|X:X::X:X) 472
clear bgp * 473
clear bgp <1-4294967295> 475
clear bgp <1-4294967295> l2vpn evpn 477
clear bgp A.B.C.D l2vpn evpn 478
clear bgp dampening 479
clear bgp external 480
clear bgp flap-statistics 482
clear bgp peer-group 483
clear bgp peer-group WORD l2vpn evpn 485
clear bgp statistics 486
clear ip bgp * (description LINE|)* 487
clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|) 488
clear ip bgp A.B.C.D 489
clear ip bgp A.B.C.D vrf 491
clear ip bgp all vrf (VRFNAME|all|default) (description LINE|) 492
clear ip bgp table-map 493
clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|) 494
debug bgp 495
distance bgp 497
exit-address-family 498
ip as-path access-list 499
ip community-list <1-99> 500
ip community-list <100-500> 501
ip community-list expanded 502
ip community-list standard 503
ip community-list WORD 504
ip extcommunity-list <1-99> 505
ip extcommunity-list <100-500> 506
ip extcommunity-list expanded 507
ip extcommunity-list standard 508
l2vpn-unnumbered-mode 509
match ip peer 510
match large-community 511
max-paths 512
max-paths eigbp 513
neighbor activate 515
neighbor advertisement-interval 516
neighbor allowas-in 517
neighbor as-origination-interval 519
neighbor attribute-unchanged 520

© 2025 IP Infusion Inc. Proprietary

neighbor authentication-key 521
neighbor capability graceful-restart 522
neighbor capability orf prefix-list 523
neighbor collide-established 525
neighbor de-activate 526
neighbor default-originate 528
neighbor description 529
neighbor disallow-infinite-holdtime 530
neighbor distribute-list 531
neighbor dont-capability-negotiate 532
neighbor ebgp-multihop 533
neighbor enforce-multihop 534
neighbor extended-optional-param 535
neighbor fall-over bfd 536
neighbor filter-list 537
neighbor g-shut 538
neighbor g-shut-timer 539
neighbor limit 540
neighbor local-as 541
neighbor maximum-prefix 543
neighbor next-hop-self 544
neighbor optional-as 545
neighbor override-capability 546
neighbor passive 547
neighbor peer-group 548
neighbor WORD peer-group range 551
neighbor port 552
neighbor prefix-list 553
neighbor remote-as 555
neighbor remove-private-AS 557
neighbor restart-time 558
neighbor route-map 559
neighbor route-reflector-client 561
neighbor route-server-client 562
neighbor send-community 563
neighbor send-community large 564
neighbor shutdown 565
neighbor strict-capability-match 567
neighbor timers 568
neighbor unsuppress-map 570
neighbor update-source 571
neighbor version 572
neighbor weight 573
neighbor WORD peer-group 575
network 576
network synchronization 578
redistribute 579

© 2025 IP Infusion Inc. Proprietary

restart bgp graceful 581
router bgp 582
set large-community 583
snmp restart bgp 584
synchronization 585
timers bgp 586
v4-unnumbered-mode 587
bgp unnumbered-mode 588

CHAPTER 2 BGP Virtual Private Network Commands 589
bgp external-route-leak 590
bgp inbound-route-filter 591
export map 592
import map 593
ip vrf 594
neighbor allow-ebgp-vpn 595
neighbor as-override 596
neighbor send-community 597
neighbor soo 598
rd (route distinguisher) 599
route-target 600

CHAPTER 3 BGP Show Commands 601
show bgp 603
show bgp A.B.C.D 604
show bgp A.B.C.D/M 606
show bgp client 607
show bgp community 608
show bgp community-list 610
show bgp dampening dampened-paths 611
show bgp dampening flap-statistics 612
show bgp dampening parameters 615
show bgp filter-list 617
show bgp inconsistent-as 619
show bgp ipv6 620
show bgp ipv6 peer-group 623
show bgp neighbors 624
show bgp neighbors advertised-routes 628
show bgp neighbors received prefix-filter 629
show bgp neighbors received-routes 630
show bgp neighbors routes 631
show bgp nexthop-tracking 633
show bgp nexthop-tree-details 635
show bgp paths 636
show bgp prefix-list 637
show bgp quote-regexp 638
show bgp regexp 639
show bgp route-map 640

© 2025 IP Infusion Inc. Proprietary

show bgp statistics . . . 641
show bgp summary . . . 643
show bgp X:X::X:X . . . 646
show bgp X:X::X:X/M longer prefixes . . . 647
show debugging bgp . . . 648
show ip bgp . . . 649
show ip bgp cidr-only . . . 652
show ip bgp community-info . . . 654
show ip bgp peer-group . . . 655
show ip bgp peer-group vrf all . . . 656
show ip bgp rtfilter all . . . 657
show ip bgp scan . . . 658
show ip extcommunity-list . . . 659
show ip protocols . . . 660
show ip vrf . . . 662
show running-config as-path access-list . . . 663
show running-config bgp . . . 664
show running-config community-list . . . 665
show running-config vrf . . . 666
show running-config ip vrf . . . 667
show running-config mac vrf . . . 668
show running-config vrf word . . . 669

Appendix D Regular Expressions . . . 670

Bidirectional Forwarding Detection Configuration . . . 671

CHAPTER 1 Base BFD Configuration . . . 672
Validation: . . . 676
Validation: . . . 678

CHAPTER 2 BFD Protocol Configurations . . . 682
OSPF—BFD Single-Hop Session . . . 682
OSPF—BFD Multi-Hop Session . . . 691
BFD Configuration in IS-IS . . . 695
BFD Configuration in BGP . . . 698

CHAPTER 3 BFD Static Route Configuration . . . 707

CHAPTER 4 BFD Authentication . . . 711
Overview . . . 711

CHAPTER 5 BFD with VRF Configuration . . . 716
Topology . . . 716
BFD Over Static Routing IPv4 and IPv6 . . . 723

CHAPTER 6 BFD Support on LAG Interface . . . 728
Overview . . . 728
Prerequisites . . . 728
Configuration . . . 729
New CLI Commands . . . 732

Glossary 732

Bidirectional Forwarding Detection Command Reference 734

CHAPTER 1 Bidirectional Forwarding Commands 735
accept-lifetime 736
bfd auth type 738
bfd 739
bfd echo 740
bfd echo interval 741
bfd echo ipv4 source 742
bfd interval 743
bfd multihop-peer 744
bfd multihop-peer A.B.C.D interval 746
bfd multihop-peer interval 748
bfd multihop-peer X:X::X:X interval 749
bfd notification 750
bfd session 751
bfd slow-timer 752
debug bfd 753
key 754
key chain 755
key-string 756
key-string encrypted 757
send-lifetime 758
show bfd 759
show bfd interface 760
show bfd session 762
show bfd session A.B.C.D 766
show bfd session ipv6 769
show debugging bfd 772
snmp restart bfd 773

CHAPTER 2 Protocol Commands for BFD 774
area virtual-link 775
bfd all-interfaces 776
debug bgp bfd 777
debug isis bfd 778
debug ospf bfd 779
ip ospf bfd 780
isis bfd 781

CHAPTER 3 BFD Static Route Commands 782
ip bfd static all-interfaces 783
ip static fall-over-bfd 784
ip static bfd 785
ipv6 bfd static all-interfaces 786
ipv6 static fall-over-bfd 787

© 2025 IP Infusion Inc. Proprietary

ipv6 static bfd 788

Intermediate System to Intermediate System Configuration 789

CHAPTER 1 IS-IS IPv4 790
Enable IS-IS on an Interface 790
Set Priority 793
Dynamic hostname 795
Redistribute Routes into IS-IS 797
Metric 803
L1-L2 Area Routing with a Single Instance 810
L1-L2 Area Routing with Multiple Instances 816
Route Leaking 822
Route Summarization 827
IS-IS Distance 832
Overload Bit 838
Passive Interface 856
IS-IS IPv4 Loop-Free Alternate Fast Reroute 861
Overview 861
Basic Configuration 862
Backup Path based on Route-Map Prefixes 869
LFA Tie-Breaker 870
LFA Termination 872
LFA For ECMP Paths 875

CHAPTER 2 IS-IS IPv6 Configuration 886
Enable IS-ISv6 on an Interface 886
Set Priority 888
Dynamic hostname 890
Redistribute Routes into IS-IS 893
Interface Metric 895
Route Summarization 899
Passive Interface 904
Enable BFD over IS-ISv6 909
Originate Default Route to ISISv6 Neighbors 912

CHAPTER 3 IS-IS Graceful Restart Configuration 915

CHAPTER 4 ISIS Multi Topology 918
Overview 918
Prerequisites 918
Configuration 921
Validation for Multi Topology 922
CLI Commands 955
Glossary 956

Intermediate System to Intermediate System Command Reference 957

CHAPTER 1 IS-IS Commands 958
accept-lifetime 961

© 2025 IP Infusion Inc. Proprietary

address-family ipv6 963
adjacency-check 964
area-password 965
authentication key-chain 966
authentication mode 967
authentication send-only 968
bfd all-interfaces 969
capability cspf 970
clear clns neighbors 971
clear clns is-neighbors 972
clear ip isis route 973
clear isis adjacency 974
clear isis counter 975
clear isis interface counter 976
clear isis process 977
debug isis 978
default-information originate 980
distance (IPv4) 981
distance (IPv6) 982
domain password 983
dynamic-hostname 984
fast-reroute per-prefix 985
fast-reroute terminate-hold-on interval 986
fast-reroute tie-break 987
ignore-lsp-errors 989
ip router isis 990
ip route high-priority tag 991
ipv6 router isis 992
isis authentication key-chain 993
isis authentication mode md5 994
isis authentication send-only 995
isis bfd 996
isis circuit-type 997
isis csnp-interval 998
isis fast-reroute per-prefix candidate disable 999
isis hello-interval 1000
isis hello-multiplier 1001
isis hello padding 1002
isis lsp-interval 1003
isis mesh-group 1004
isis metric 1005
isis network 1006
isis password 1007
isis priority 1008
isis retransmit-interval 1009
ispf 1010
isis wait-timer 1011

© 2025 IP Infusion Inc. Proprietary

isis wide-metric 1012
isis tag 1013
isis te-metric 1014
is-type 1015
key chain 1016
key 1017
key-string 1018
key-string encrypted 1019
lsp-gen-interval 1020
lsp-mtu 1021
lsp-refresh-interval 1022
max-area-address 1023
max-lsp-lifetime 1024
metric-style 1025
net 1027
passive-interface 1028
prc-interval-exp 1029
redistribute 1030
redistribute isis 1032
redistribute isis WORD 1033
router isis 1034
send-lifetime 1035
set-overload-bit 1037
snmp restart isis 1038
spf-interval-exp 1039
summary-address 1040
summary-prefix 1041

CHAPTER 2 IS-IS Graceful Restart Commands 1042
capability restart graceful 1043
isis restart grace-period 1044
isis restart-hello-interval 1045
isis restart helper 1046
isis restart suppress-adjacency 1047
restart isis graceful 1048
restart-timer 1049

CHAPTER 3 IS-IS Show Commands 1050
show clns is-neighbors 1051
show clns neighbors 1053
show debugging isis 1055
show ip isis igp-shortcut-lsp 1056
show ip isis route 1057
show ip isis route igp-shortcut 1059
show ip protocols 1061
show ip route fast-reroute 1062
show ip isis route fast-reroute 1063
show ip isis lfa-config 1064

show isis counter 1065
show isis database 1066
show isis interface 1069
show isis spf-logs 1072
show isis topology 1073
show running-config interface isis 1075
show running-config router isis 1076

Open Shortest Path First Configuration 1077

CHAPTER 1 OSPFv2 1078
Enable OSPF on an Interface 1078
Set Priority 1082
Area Border Router 1087
Redistribute Routes into OSPF 1091
Cost 1093
Path MTU for IPv4 and IPv6 1098
Virtual Links 1101
OSPF Authentication 1105
Multiple OSPF Instances 1109
Multiple OSPF Instances on Same Subnet 1120
Multi-Area Adjacency Configuration 1123
LSA Throttling 1126
SNMP Support for Multiple Instance of OSPF Configuration 1149
OSPF Stub Router Advertisement 1171

CHAPTER 2 Multi-Area Redundant Adjacency Configuration 1179
Overview 1179
Prerequisites 1180
Configuration 1180
Implementation Example 1193
Revised CLI Commands 1193
Glossary 1194

CHAPTER 3 OSPFv3 1196
Enable OSPFv3 on an Interface 1196
Set Priority 1199
Area Border Router 1203
Redistribute Routes into OSPFv3 1209
Cost 1215
Virtual Links 1225
Multiple Instances 1232
Not-So-Stubby Area 1243
NSSA with the Summary Address Option 1247
NSSA with the Translator Role Option 1251
Link LSA Suppression 1255
Originate Type-7 LSAs and Translate to Type-5 1257
Summarize Inter-Area and External Routes 1262

© 2025 IP Infusion Inc. Proprietary

Distribute List 1267
OSPFv3 Authentication 1271

CHAPTER 4 OSPFv3 Graceful Restart Configuration 1278

Open Shortest Path First Command Reference 1285

CHAPTER 1 OSPFv2 Commands 1286
area authentication 1289
area default-cost 1290
area interface 1291
area interface authentication 1292
area interface network-type 1293
area interface passive 1294
area filter-list 1295
area nssa 1296
area range 1298
area shortcut 1299
area stub 1300
area virtual-link 1301
auto-cost reference bandwidth 1304
bfd all-interfaces 1305
capability cspf 1306
capability lls 1307
capability opaque 1308
capability restart 1309
capability traffic-engineering 1310
capability vrf-lite 1311
clear ip ospf 1312
compatible rfc1583 1313
debug ospf 1314
debug ospf database-timer rate-limit 1315
debug ospf events 1316
debug ospf ifsm 1317
debug ip ospf graceful-restart 1318
debug ip ospf lfa 1319
debug ip ospf redist 1320
debug ip ospf retransmission 1321
debug ospf lsa 1322
debug ospf nfsm 1323
debug ospf nsm 1324
debug ospf packet 1325
debug ospf rib 1326
debug ospf route 1327
default-information originate 1328
default-metric 1330
distance 1331
distribute-list 1332

© 2025 IP Infusion Inc. Proprietary

enable db-summary-opt 1334
fast-reroute keep-all-paths 1335
fast-reroute tie-break 1336
host area 1338
ip ospf authentication 1339
ip ospf authentication-key 1340
ip ospf bfd 1342
ip ospf cost 1343
ip ospf database-filter 1344
ip ospf dead-interval 1345
ip ospf disable 1346
ip ospf fast-reroute per-prefix candidate disable 1347
ip ospf flood-reduction 1348
ip ospf hello-interval 1349
ip ospf multi-area 1350
ip ospf message-digest-key 1351
ip ospf mtu 1353
ip ospf mtu-ignore 1354
ip ospf network 1355
ip ospf priority 1356
ip ospf retransmit-interval 1357
ip ospf transmit-delay 1358
log-adjacency-changes 1359
max-concurrent-dd 1360
maximum-area 1361
max-metric 1362
neighbor 1364
network 1365
ospf abr-type 1367
ospf area-interface-config-mode 1368
ospf flood-reduction 1369
ospf restart grace-period 1370
ospf restart helper 1371
ospf router-id 1372
overflow database 1373
overflow database external 1374
passive-interface 1375
redistribute 1376
restart ip ospf graceful 1378
router ospf 1379
show debugging ospf 1380
show ip ospf 1381
show ip ospf border-routers 1385
show ip ospf database brief 1386
show ip ospf database detail 1388
show ip ospf igp-shortcut-lsp 1395
show ip ospf igp-shortcut-route 1396

© 2025 IP Infusion Inc. Proprietary

show ip ospf interface 1397
show ip ospf multi-area-adjacencies 1400
show ip ospf neighbor 1402
show ip ospf route 1406
show ip ospf valid 1408
show ip ospf virtual-links 1409
show ip protocols 1411
show ip route fast-reroute 1413
shutdown 1414
snmp context-name 1415
snmp restart ospf 1416
summary-address 1417
timers lsa arrival 1418
timers spf exp 1419
timers throttle lsa 1420

CHAPTER 2 OSPFv2 Graceful Restart Commands 1421
action interface 1422
capability restart 1423
event nbr 1424
if-event-track ospf event 1425
debug ip ospf graceful-restart 1426
ospf restart grace-period 1427
ospf restart helper 1428
restart ospf graceful 1429

CHAPTER 3 OSPFv3 Commands 1430
abr-type 1432
area default-cost 1433
area nssa 1434
area range 1436
area stub 1437
area virtual-link 1438
auto-cost reference bandwidth 1440
bfd all-interfaces 1441
capability restart 1442
clear ipv6 ospf process 1443
debug ipv6 ospf 1444
debug ipv6 ospf bfd 1445
debug ipv6 ospf events 1446
debug ipv6 ospf ifsm 1447
debug ipv6 ospf lfa 1448
debug ipv6 ospf lsa 1449
debug ipv6 ospf nfsm 1450
debug ipv6 ospf nsm 1451
debug ipv6 ospf packet 1452
debug ipv6 ospf retransmission 1453
debug ipv6 ospf rib 1454

© 2025 IP Infusion Inc. Proprietary

debug ipv6 ospf route 1455
default-information originate 1456
default-metric 1458
distance 1459
distribute-list 1460
enable db-summary-opt 1462
fast-reroute keep-all-paths 1463
fast-reroute tie-break 1464
ipv6 ospf authentication 1466
ipv6 ospf bfd 1467
ipv6 ospf cost 1468
ipv6 ospf dead-interval 1469
ipv6 ospf display route single-line 1470
ipv6 ospf hello-interval 1471
ipv6 ospf link-lsa-suppression 1472
ipv6 ospf mtu-ignore 1473
ipv6 ospf neighbor 1474
ipv6 ospf network 1476
ipv6 ospf priority 1477
ipv6 ospf restart grace-period 1478
ipv6 ospf restart helper 1479
ipv6 ospf restart planned-only 1480
ipv6 ospf retransmit-interval 1481
ipv6 ospf transmit-delay 1482
ipv6 router ospf 1483
ipv6 te-metric 1485
log-adjacency-changes 1486
max-concurrent-dd 1487
passive-interface 1488
redistribute 1489
restart ipv6 ospf graceful 1491
router-id 1492
router ipv6 ospf 1493
show debugging ipv6 ospf 1494
show ipv6 ospf 1495
show ipv6 ospf database 1497
show ipv6 ospf interface 1501
show ipv6 ospf neighbor 1503
show ipv6 ospf route 1506
show ipv6 route fast-reroute 1508
show ipv6 ospfv3 topology 1509
show ipv6 ospf virtual-links 1511
show ipv6 vrf 1513
shutdown 1514
snmp restart ospf6 1515
summary-address 1516
timers spf exp 1518

© 2025 IP Infusion Inc. Proprietary

CHAPTER 4 OSPFv3 Graceful Restart Commands 1519
capability restart 1520
ipv6 ospf restart grace-period 1521
ipv6 ospf restart helper 1522
ipv6 ospf restart planned-only 1523
restart ipv6 ospf graceful 1524

CHAPTER 5 OSPF VPN Commands 1525
capability vrf-lite 1526
router ospf vrf 1527
domain-id 1528

Routing Information Protocol Configuration 1530

CHAPTER 1 RIP 1531
Enable RIP 1531
Specify RIP Version 1535
Authentication with a Single Key 1538
Text Authentication with Multiple Keys 1543
MD5 Authentication with Multiple Keys 1550

CHAPTER 2 RIPng 1558
Configuration 1558

Routing Information Protocol Command Reference 1559

CHAPTER 1 Routing Information Protocol Commands 1560
accept-lifetime 1562
cisco-metric-behavior 1564
clear ip rip route 1565
clear ip rip route vrf NAME 1566
clear ip rip statistics 1567
debug rip 1568
default-information originate 1569
default-metric 1570
distance 1571
distribute-list 1572
ip rip authentication key-chain 1573
ip rip authentication mode 1574
ip rip authentication string 1575
ip rip receive-packet 1576
ip rip receive version 1577
ip rip send-packet 1578
ip rip send version 1579
ip rip split-horizon 1580
key 1581
key chain 1582
key-string 1583
maximum-prefix 1584

neighbor . . . 1585
network . . . 1586
offset-list . . . 1587
passive-interface . . . 1588
recv-buffer-size . . . 1589
redistribute . . . 1590
route . . . 1592
router rip . . . 1593
send-lifetime . . . 1594
show debugging rip . . . 1596
show ip protocols rip . . . 1597
show ip rip . . . 1599
show ip rip interface . . . 1601
show ip rip statistics . . . 1603
snmp restart rip . . . 1605
timers basic . . . 1606
version . . . 1607

CHAPTER 2 RIPng Commands . . . 1608
aggregate-address . . . 1609
cisco-metric-behavior . . . 1610
clear ipv6 rip route . . . 1611
debug ipv6 rip . . . 1612
default-information originate . . . 1613
default-metric . . . 1614
distance . . . 1615
distribute-list . . . 1616
ipv6 rip metric-offset . . . 1617
ipv6 rip split-horizon . . . 1618
ipv6 router rip . . . 1619
neighbor . . . 1620
offset-list . . . 1621
passive-interface . . . 1622
recv-buffer-size . . . 1623
redistribute . . . 1624
route . . . 1625
route-map . . . 1626
router ipv6 rip . . . 1627
show debugging ipv6 rip . . . 1628
show ipv6 protocols rip . . . 1629
show ipv6 rip . . . 1630
show ipv6 rip interface . . . 1631
timers basic . . . 1632

CHAPTER 3 Routing Information Protocol VPN Commands . . . 1633
show ip rip interface vrf . . . 1634
show ip rip vrf . . . 1636
show ip vrf . . . 1638

© 2025 IP Infusion Inc. Proprietary

CHAPTER 4 Routing Information Protocol Authentication 1639
Single Key Authentication 1639
Multiple Keys Authentication. 1639

VRF Lite Configuration 1641

CHAPTER 1 VRF Configuration. 1642
Overview 1642
Default VRF. 1642
User-Defined VRF 1642

CHAPTER 2 OSPF Configuration 1644
Overview 1644
Configuration IPv4 VRF. 1644

CHAPTER 3 RIP Configuration 1646
Configuration IPv4 VRF. 1646
Configuration IPv6 VRF. 1648

CHAPTER 4 ISIS Configuration 1651
Overview 1651
Topology 1651
Configuration IPv4 VRF. 1651

CHAPTER 5 BGP Configuration 1654
Overview 1654
Configuration. 1654

CHAPTER 6 Inter-VRF Route Leaking Configuration 1658
Overview 1658
Static Leaking 1658
Dynamic Leaking 1660

Virtual Router Redundancy Protocol Configuration 1665

CHAPTER 1 VRRP Configuration 1666
Terminology 1666
VRRP Process 1666
One Virtual Router 1668
Two Virtual Routers 1671
Two Backup Routers 1675
Interface Tracking. 1680
VRRP-Backward Compatibility. 1687
Redundancy Using VRRP and OSPF: Two Virtual Routers 1690
VRRP Over MLAG 1695
OBJECT Tracking Using IPSLA. 1702
Custom VRF Configuration. 1709

CHAPTER 2 Object Tracking Using IP SLA. 1716
Overview 1716
Topology 1716

Validation 1720

CHAPTER 3 VRRP IPv6 Configuration 1723
Configuration 1723
Disabling the Master 1726
GLOBAL IPV6 SUPPORT FOR VRRP 1727

Virtual Router Redundancy Protocol Command Reference 1731

CHAPTER 1 VRRP Commands 1732
accept-mode 1733
advertisement-interval 1734
authentication text 1735
circuit-failover 1736
debug vrrp 1737
disable 1738
enable 1739
operational-ip 1740
preempt-mode 1741
priority 1742
router vrrp 1743
show debugging vrrp 1744
show running-config router vrrp 1745
show vrrp 1746
show vrrp summary 1748
show vrrp <1-255> 1749
show vrrp (global | ipv4) statistics 1751
snmp restart vrrp 1752
switch-back-delay 1753
virtual-ip 1754
vrrp compatible-v2 1755
vrrp vmac 1756

CHAPTER 2 VRRP v6 Commands 1757
advertisement-interval 1758
circuit-failover 1759
disable 1760
enable 1761
preempt-mode 1762
priority 1763
router ipv6 vrrp 1764
router ipv6 vrrp vlan 1765
show running-config router ipv6 vrrp 1766
virtual-ipv6 1767

Index 1768

© 2025 IP Infusion Inc. Proprietary

Preface

This guide describes how to configure OcNOS.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions

Table 1 on page 26 shows the conventions used in this guide.

Table 1: Conventions

Convention	Description
Italics	Emphasized terms; titles of books
Note:	Special instructions, suggestions, or warnings
`monospaced type`	Code elements such as commands, parameters, files, and directories

Chapter Organization

The chapters in command references are organized as described in Command Description Format.

The chapters in configuration guides are organized into these major sections:

- An overview that explains a configuration in words
- Topology with a diagram that shows the devices and connections used in the configuration
- Configuration steps in a table for each device where the left-hand side shows the commands you enter and the right-hand side explains the actions that the commands perform
- Validation which shows commands and their output that verify the configuration

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

The features described in this document that are available depend upon the OcNOS SKU that you purchased. See the *Feature Matrix* for a description of the OcNOS SKUs.

Migration Guide

Check the *Migration Guide* for configuration changes to make when migrating from one version of OcNOS to another.

Support

For support-related questions, contact support@ipinfusion.com.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

Command Line Interface

This chapter introduces the OcNOS Command Line Interface (CLI) and how to use its features.

Overview

You use the CLI to configure, monitor, and maintain OcNOS devices. The CLI is text-based and each command is usually associated with a specific task.

You can give the commands described in this manual locally from the console of a device running OcNOS or remotely from a terminal emulator such as `putty` or `xterm`. You can also use the commands in scripts to automate configuration tasks.

Command Line Interface Help

You access the CLI help by entering a full or partial command string and a question mark “?”. The CLI displays the command keywords or parameters along with a short description. For example, at the CLI command prompt, type:

```
> show ?
```

The CLI displays this keyword list with short descriptions for each keyword:

```
show ?
  application-priority            Application Priority
  arp                             Internet Protocol (IP)
  bfd                             Bidirectional Forwarding Detection (BFD)
  bgp                             Border Gateway Protocol (BGP)
  bi-lsp                          Bi-directional lsp status and configuration
  bridge                          Bridge group commands
  ce-vlan                         COS Preservation for Customer Edge VLAN
  class-map                       Class map entry
  cli                             Show CLI tree of current mode
  clns                            Connectionless-Mode Network Service (CLNS)
  control-adjacency               Control Adjacency status and configuration
  control-channel                 Control Channel status and configuration
  cspf                            CSPF Information
  customer                        Display Customer spanning-tree
  cvlan                           Display CVLAN information
  debugging                       Debugging functions
  etherchannel                    LACP etherchannel
  ethernet                        Layer-2
...
```

If you type the ? in the middle of a keyword, the CLI displays help for that keyword only.

```
> show de?
debugging  Debugging functions
```

If you type the ? in the middle of a keyword, but the incomplete keyword matches several other keywords, OcNOS displays help for all matching keywords.

```
> show i? (CLI does not display the question mark).
interface  Interface status and configuration
ip         IP information
isis       ISIS information
```

Command Completion

The CLI can complete the spelling of a command or a parameter. Begin typing the command or parameter and then press the tab key. For example, at the CLI command prompt type `sh`:

```
> sh
```

Press the tab key. The CLI displays:

```
> show
```

If the spelling of a command or parameter is ambiguous, the CLI displays the choices that match the abbreviation. Type `show i` and press the tab key. The CLI displays:

```
> show i
  interface  ip          ipv6       isis
> show i
```

The CLI displays the `interface` and `ip` keywords. Type `n` to select `interface` and press the tab key. The CLI displays:

```
> show in
> show interface
```

Type ? and the CLI displays the list of parameters for the `show interface` command.

```
> show interface
  IFNAME  Interface name
  |       Output modifiers
  >       Output redirection
  <cr>
```

The CLI displays the only parameter associated with this command, the `IFNAME` parameter.

Command Abbreviations

The CLI accepts abbreviations that uniquely identify a keyword in commands. For example:

```
> sh int xe0
```

is an abbreviation for:

```
> show interface xe0
```

Command Line Errors

Any unknown spelling causes the CLI to display the error `Unrecognized command` in response to the ?. The CLI displays the command again as last entered.

```
> show dd?
% Unrecognized command
> show dd
```

When you press the Enter key after typing an invalid command, the CLI displays:

```
(config)#router ospf here
                        ^
% Invalid input detected at '^' marker.
```

where the ^ points to the first character in error in the command.

If a command is incomplete, the CLI displays the following message:

```
> show
% Incomplete command.
```

Some commands are too long for the display line and can wrap mid-parameter or mid-keyword, as shown below. This does *not* cause an error and the command performs as expected:

```
area 10.10.0.18 virtual-link 10.10.0.19 authent
ication-key 57393
```

Command Negation

Many commands have a `no` form that resets a feature to its default value or disables the feature. For example:

- The `ip address` command assigns an IPv4 address to an interface
- The `no ip address` command removes an IPv4 address from an interface

Syntax Conventions

Table 2 on page 30 describes the conventions used to represent command syntax in this reference.

Table 2: Syntax conventions

Convention	Description	Example
`monospaced font`	Command strings entered on a command line	`show ip ospf`
`lowercase`	Keywords that you enter exactly as shown in the command syntax.	`show ip ospf`
`UPPERCASE`	See Variable Placeholders	`IFNAME`
`()`	Optional parameters, from which you must select one. Vertical bars delimit the selections. Do not enter the parentheses or vertical bars as part of the command.	`(A.B.C.D|<0-4294967295>)`
`()`	Optional parameters, from which you select one or none. Vertical bars delimit the selections. Do not enter the parentheses or vertical bars as part of the command.	`(A.B.C.D|<0-4294967295>|)`
`()`	Optional parameter which you can specify or omit. Do not enter the parentheses or vertical bar as part of the command.	`(IFNAME|)`
`{}`	Optional parameters, from which you must select one or more. Vertical bars delimit the selections. Do not enter the braces or vertical bars as part of the command.	`{intra-area <1-255>|inter-area <1-255>|external <1-255>}`

© 2025 IP Infusion Inc. Proprietary

Table 2: Syntax conventions (Continued)

Convention	Description	Example
`[]`	Optional parameters, from which you select zero or more. Vertical bars delimit the selections. Do not enter the brackets or vertical bars as part of the command.	`[<1-65535>|AA:NN|internet|local-AS|no-advertise|no-export]`
`?`	Nonrepeatable parameter. The parameter that follows a question mark can only appear once in a command string. Do not enter the question mark as part of the command.	`?route-map WORD`
`.`	Repeatable parameter. The parameter that follows a period can be repeated more than once. Do not enter the period as part of the command.	`set as-path prepend .<1-65535>`

Variable Placeholders

Table 3 on page 31 shows the tokens used in command syntax use to represent variables for which you supply a value.

Table 3: Variable placeholders

Token	Description
`WORD`	A contiguous text string (excluding spaces)
`LINE`	A text string, including spaces; no other parameters can follow this parameter
`IFNAME`	Interface name whose format varies depending on the platform; examples are: `eth0`, `Ethernet0`, `ethernet0`, `xe0`
`A.B.C.D`	IPv4 address
`A.B.C.D/M`	IPv4 address and mask/prefix
`X:X::X:X`	IPv6 address
`X:X::X:X/M`	IPv6 address and mask/prefix
`HH:MM:SS`	Time format
`AA:NN`	BGP community value
`XX:XX:XX:XX:XX:XX`	MAC address

`<1-5>`
`<1-65535>`
`<0-2147483647>`
`<0-4294967295>`
	Numeric range

Command Description Format

Table 4 on page 32 explains the sections used to describe each command in this reference.

Table 4: Command descriptions

Section	Description
Command Name	The name of the command, followed by what the command does and when should it be used
Command Syntax	The syntax of the command
Parameters	Parameters and options for the command
Default	The state before the command is executed
Command Mode	The mode in which the command runs; see Command Modes
Example	An example of the command being executed

Keyboard Operations

Table 5 on page 32 lists the operations you can perform from the keyboard.

Table 5: Keyboard operations

Key combination	Operation
Left arrow or Ctrl+b	Moves one character to the left. When a command extends beyond a single line, you can press left arrow or Ctrl+b repeatedly to scroll toward the beginning of the line, or you can press Ctrl+a to go directly to the beginning of the line.
Right arrow or Ctrl-f	Moves one character to the right. When a command extends beyond a single line, you can press right arrow or Ctrl+f repeatedly to scroll toward the end of the line, or you can press Ctrl+e to go directly to the end of the line.
Esc, b	Moves back one word
Esc, f	Moves forward one word
Ctrl+e	Moves to end of the line
Ctrl+a	Moves to the beginning of the line
Ctrl+u	Deletes the line
Ctrl+w	Deletes from the cursor to the previous whitespace
Alt+d	Deletes the current word
Ctrl+k	Deletes from the cursor to the end of line
Ctrl+y	Pastes text previously deleted with Ctrl+k, Alt+d, Ctrl+w, or Ctrl+u at the cursor

© 2025 IP Infusion Inc. Proprietary

Table 5: Keyboard operations (Continued)

Key combination	Operation
Ctrl+t	Transposes the current character with the previous character
Ctrl+c	Ignores the current line and redisplays the command prompt
Ctrl+z	Ends configuration mode and returns to exec mode
Ctrl+l	Clears the screen
Up Arrow or Ctrl+p	Scroll backward through command history
Down Arrow or Ctrl+n	Scroll forward through command history

Show Command Modifiers

You can use two tokens to modify the output of a `show` command. Enter a question mark to display these tokens:

```
# show users ?
  | Output modifiers
  > Output redirection
```

You can type the | (vertical bar character) to use output modifiers. For example:

```
> show rsvp | ?
  begin     Begin with the line that matches
  exclude   Exclude lines that match
  include   Include lines that match
  last      Last few lines
  redirect  Redirect output
```

Begin Modifier

The `begin` modifier displays the output beginning with the first line that contains the input string (everything typed after the `begin` keyword). For example:

```
# show running-config | begin xe1
...skipping
interface xe1
 ipv6 address fe80::204:75ff:fee6:5393/64
!
interface xe2
 ipv6 address fe80::20d:56ff:fe96:725a/64
!
line con 0
 login
!
end
```

You can specify a regular expression after the `begin` keyword, This example begins the output at a line with either “xe2” or “xe4”:

```
# show running-config | begin xe[2-4]

...skipping
```

```
interface xe2
 shutdown
!
interface xe4
 shutdown
!
interface svlan0.1
 no shutdown
!
route-map myroute permit 2
!
route-map mymap1 permit 10
!
route-map rmap1 permit 2
!
line con 0
 login
line vty 0 4
 login
!
end
```

Include Modifier

The `include` modifier includes only those lines of output that contain the input string. In the output below, all lines containing the word “input” are included:

```
# show interface xe1 | include input
    input packets 80434552, bytes 2147483647, dropped 0, multicast packets 0
    input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 1, missed 0
```

You can specify a regular expression after the `include` keyword. This examples includes all lines with “input” or “output”:

```
#show interface xe0 | include (in|out)put
    input packets 597058, bytes 338081476, dropped 0, multicast packets 0
    input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 0, missed 0
    output packets 613147, bytes 126055987, dropped 0
    output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
```

Exclude Modifier

The `exclude` modifier excludes all lines of output that contain the input string. In the following output example, all lines containing the word “input” are excluded:

```
# show interface xe1 | exclude input
Interface xe1
  Scope: both
  Hardware is Ethernet, address is 0004.75e6.5393
  index 3 metric 1 mtu 1500 <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Administrative Group(s): None
  DSTE Bandwidth Constraint Mode is MAM
  inet6 fe80::204:75ff:fee6:5393/64
    output packets 4438, bytes 394940, dropped 0
    output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
    collisions 0
```

You can specify a regular expression after the `exclude` keyword. This example excludes lines with “output” or “input”:

```
# show interface xe0 | exclude (in|out)put
Interface xe0
  Scope: both
  Hardware is Ethernet  Current HW addr: 001b.2139.6c4a
  Physical:001b.2139.6c4a  Logical:(not set)
  index 2 metric 1 mtu 1500 duplex-full arp ageing timeout 3000
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Bandwidth 100m
  DHCP client is disabled.
  inet 10.1.2.173/24 broadcast 10.1.2.255
  VRRP Master of :  VRRP is not configured on this interface.
  inet6 fe80::21b:21ff:fe39:6c4a/64
    collisions 0
```

Redirect Modifier

The `redirect` modifier writes the output into a file. The output is not displayed.

```
# show cli history | redirect /var/frame.txt
```

The output redirection token (>) does the same thing:

```
# show cli history >/var/frame.txt
```

Last Modifier

The `last` modifier displays the output of last few number of lines (As per the user input). The last number ranges from 1 to 9999.

For example:

```
#show running-config | last 10
```

String Parameters

The restrictions in Table 6 on page 36 apply for all string parameters used in OcNOS commands, unless some other restrictions are noted for a particular command.

Table 6: String parameter restrictions

Restriction	Description
Input length	1965 characters or less
Restricted special characters	“?”, “,”, “>”, “|”, and “=” The “|” character is allowed only for the `description` command in interface mode.

Command Modes

Commands are grouped into modes arranged in a hierarchy. Each mode has its own set of commands. Table P-7 lists the command modes common to all protocols.

Table 7: Common command modes

Name	Description
Executive mode	Also called *view* mode, this is the first mode to appear after you start the CLI. It is a base mode from where you can perform basic commands such as `show`, `exit`, `quit`, `help`, and `enable`.
Privileged executive mode	Also called *enable* mode, in this mode you can run additional basic commands such as `debug`, `write`, and `show`.
Configure mode	Also called *configure terminal* mode, in this mode you can run configuration commands and go into other modes such as interface, router, route map, key chain, and address family. Configure mode is single user. Only one user at a time can be in configure mode.
Interface mode	In this mode you can configure protocol-specific settings for a particular interface. Any setting you configure in this mode overrides a setting configured in router mode.
Router mode	This mode is used to configure router-specific settings for a protocol such as BGP or OSPF.

Command Mode Tree

The diagram below shows the common command mode hierarchy.

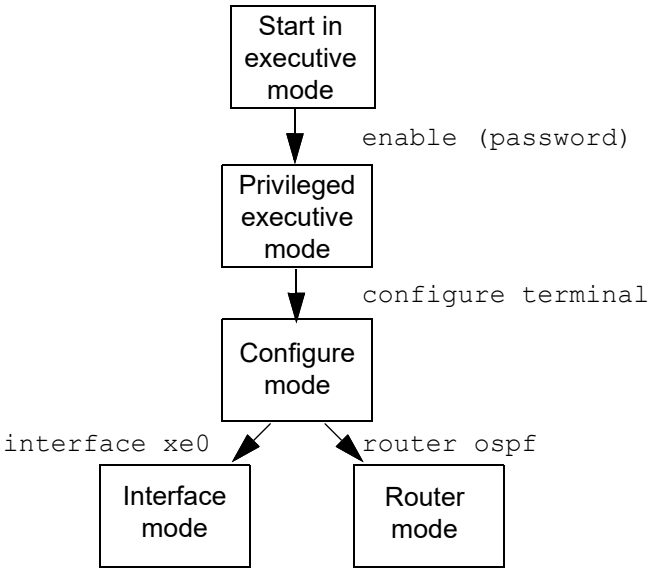

Figure P-1: Common command modes

To change modes:

1. Enter privileged executive mode by entering `enable` in Executive mode.
2. Enter configure mode by entering `configure terminal` in Privileged Executive mode.

The example below shows moving from executive mode to privileged executive mode to configure mode and finally to router mode:

```
> enable mypassword
# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)# router ospf
(config-router)#
```

Note: Each protocol can have modes in addition to the common command modes. See the command reference for the respective protocol for details.

Transaction-based Command-line Interface

The OcNOS command line interface is transaction based:

- Any changes done in configure mode are stored in a separate *candidate* configuration that you can view with the `show transaction current` command.
- When a configuration is complete, apply the candidate configuration to the running configuration with the `commit` command.
- If a `commit` fails, no configuration is applied as the entire transaction is considered failed. You can continue to change the candidate configuration and then retry the `commit`.
- Discard the candidate configuration with the `abort transaction` command.
- Check the last aborted transaction with the `show transaction last-aborted` command.
- Multiple configurations cannot be removed with a single `commit` . You must remove each configuration followed by a `commit`

Note: All commands MUST be executed only in the default CML shell (`cmlsh`). If you log in as root and start `imish`, then the system configurations will go out of sync. The `imish` shell is not supported and should not be started manually.

Unicast Configuration

CHAPTER 1 Static Routes

This chapter contains basic static routing configuration examples.

This example shows the complete configuration to enable static routing in a simple network topology. A static route is composed of a network prefix (host address) and a nexthop (gateway). Static routes are useful in small networks. They are simple solutions for making a few destinations reachable. Large networks use dynamic routing protocols.

Topology

Router R1 is configured with these static routes:

- The remote network 10.10.12.0/24
- The loopback address (host addresses) of router R2
- The loopback address of router R3

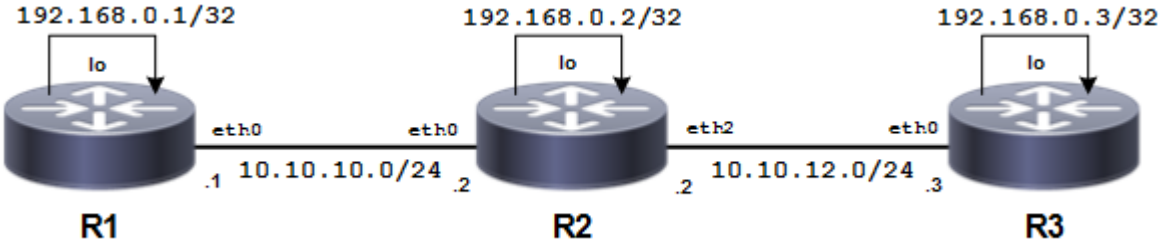

Figure 1-1: Basic Static Route

In all three routes, interface `eth0` of router R2 is the gateway. Router R3 is configured with a default static route that is equivalent to configuring separate static routes with the same gateway or nexthop address. Router R2 has two routes, one for each of the remote routers' loopback address.

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 192.168.0.1/32 secondary`	Configure the IP address on this interface, and specify a 32-bit mask, making it a host address.
`(config-if)#exit`	Exit interface mode.

`(config)#ip route 10.10.12.0/24 10.10.10.2`
`(config)#ip route 192.168.0.2/32 10.10.10.2`
`(config)#ip route 192.168.0.3/32 10.10.10.2`
	Specify the destination prefix and mask for the network and a gateway.
	Because R2 is the only next hop available, you can configure a default route instead of configuring the same static route for individual addresses. See the configuration of R3.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 192.168.0.2/32`	Configure the IP address on this interface, and specify a 32-bit mask, making it a host address.
`(config-if)#exit`	Exit interface mode.
`(config)#ip route 192.168.0.1/32 10.10.10.1` `(config)#ip route 192.168.0.3/32 10.10.12.3`	Specify the destination and mask for the network and a gateway.

R3

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 192.168.0.3/32`	Configure the IP address on this interface, and specify a 32-bit mask, making it a host address.
`(config-if)#exit`	Exit interface mode.
`(config)#ip route 0.0.0.0/0 10.10.12.2`	Specify 10.10.12.2 as a default gateway to reach any network. Because 10.10.12.2 is the only available route, you can specify it as the default gateway instead of specifying it as the gateway for an individual network or host address.

Validation

show ip route, show ip route summary, show ip route database

R1

```
#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.4.1 to network 0.0.0.0

K*      0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C       10.10.10.0/24 is directly connected, eth1
S       10.10.12.0/24 [1/0] via 10.10.10.2, eth1
C       10.12.4.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo
C       192.168.0.1/32 is directly connected, lo
S       192.168.0.2/32 [1/0] via 10.10.10.2, eth1
S       192.168.0.3/32 [1/0] via 10.10.10.2, eth1

#show ip route summary
IP routing table name is Default-IP-Routing-Table(0)
```

```
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 8
Total number of IPv4 paths       : 8
Route Source    Networks
kernel          1
connected       4
static          3
Total           8
FIB             0

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv4 ECMP routes   : 0
 Total number of IPv4 ECMP paths    : 0

---------------------------------

#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
K    *> 0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C    *> 10.10.10.0/24 is directly connected, eth1
S    *> 10.10.12.0/24 [1/0] via 10.10.10.2, eth1
C    *> 10.12.4.0/24 is directly connected, eth0
C    *> 127.0.0.0/8 is directly connected, lo
C    *> 192.168.0.1/32 is directly connected, lo
S    *> 192.168.0.2/32 [1/0] via 10.10.10.2, eth1
S    *> 192.168.0.3/32 [1/0] via 10.10.10.2, eth1

Gateway of last resort is not set
```

R2

```
#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.4.1 to network 0.0.0.0

K*      0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C       10.10.10.0/24 is directly connected, eth1
C       10.10.12.0/24 is directly connected, eth2
C       10.12.4.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo
S       192.168.0.1/32 [1/0] via 10.10.10.1, eth1
```

```
C       192.168.0.2/32 is directly connected, lo
S       192.168.0.3/32 [1/0] via 10.10.12.3, eth2

#sh ip route summary
IP routing table name is Default-IP-Routing-Table(0)
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 9
Total number of IPv4 paths       : 9
Route Source    Networks
kernel          1
connected       5
static          3
Total           9
FIB             0

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv4 ECMP routes   : 0
 Total number of IPv4 ECMP paths    : 0

---------------------------------

#sh ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
K    *> 0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C    *> 10.10.10.0/24 is directly connected, eth1
C    *> 10.10.12.0/24 is directly connected, eth2
C    *> 10.12.4.0/24 is directly connected, eth0
C    *> 127.0.0.0/8 is directly connected, lo
S    *> 192.168.0.1/32 [1/0] via 10.10.10.1, eth1
C    *> 192.168.0.2/32 is directly connected, lo
S    *> 192.168.0.3/32 [1/0] via 10.10.12.3, eth2

Gateway of last resort is not set
```

R3

```
#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.4.1 to network 0.0.0.0

K*      0.0.0.0/0 [0/0] via 10.12.4.1, eth0
```

```
C       10.10.12.0/24 is directly connected, eth2
C       10.12.4.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo
C       192.168.0.3/32 is directly connected, lo

#sh ip route summary
IP routing table name is Default-IP-Routing-Table(0)
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 6
Total number of IPv4 paths       : 6
Route Source    Networks
kernel          2
connected       4
Total           6
FIB             0

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv4 ECMP routes   : 0
 Total number of IPv4 ECMP paths    : 0
---------------------------------


#sh ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
K    *> 0.0.0.0/0 [0/0] via 10.12.4.1, eth0
S       0.0.0.0/0 [1/0] via 10.10.12.2 inactive
C    *> 10.10.12.0/24 is directly connected, eth2
C    *> 10.12.4.0/24 is directly connected, eth0
C    *> 127.0.0.0/8 is directly connected, lo
C    *> 192.168.0.3/32 is directly connected, lo

Gateway of last resort is not set
```

CHAPTER 2 Static Route Discard Configuration

This chapter show how to configure the static route discard feature.

Overview

If you identify some routers/attackers distributing invalid/bogus routes just to use the resources of the device or to make the device unstable, you can configure route-map rules, discard all routes, and black hole traffic corresponding to those routes.

To so this, you add “discard” route entries for a prefix in a route map with the “set interface null0” command. You then apply that route map to a BGP neighbor.

IPv4 Route Discard

Figure 2-2 shows the configuration required to enable static route discard for IPv4.

Topology

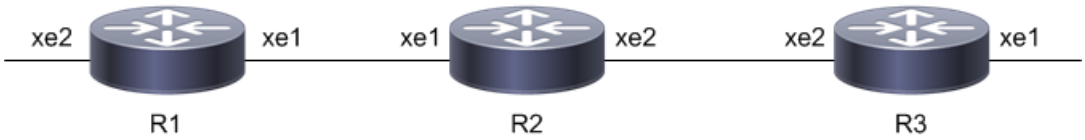

Figure 2-2: Static route discard topology

Configuration

R1

`R1#configure terminal`	Enter configure mode
`R1(config)#interface xe1`	Enter interface mode for xe1
`R1(config-if)#ip address 2.2.2.2/24`	Assign an IP address to the interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface xe2`	Enter interface mode for xe1
`R1(config-if)#ip address 1.1.1.2/24`	Assign an IP address to the interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode
`R1(config)#router bgp 2`	Enter BGP router mode
`R1(config-router)#neighbor 2.2.2.3 remote-as 3`	Create static BGP neighbor 2.2.2.3 with remote autonomous system value 3
`R1(config-router)#address-family ipv4`	Enter address-family mode

`R1(config-router-af)#neighbor 2.2.2.3 activate`	Activate neighbor 2.2.2.3
`R1(config-router-af)#redistribute connected`	Advertise the connected network into BGP
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-router-af)#end`	Exit BGP router mode

R2

`R2#configure terminal`	Enter configure mode
`R2(config)#ip prefix-list p1`	Configure IP prefix list p1
`R2(config-ip-prefix-list)#permit any`	Configure permit any
`R2(config)#route-map r1`	Enter route-map mode
`R2(config-route-map)#match ip address prefix-list p1`	Configure match ip prefix list p1
`R2(config-route-map)#set interface null0`	Set the interface to null0
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-route-map)#exit`	Exit route-map mode.
`R2(config)#interface xe1`	Enter interface mode for xe1
`R2(config-if)#ip address 2.2.2.3/24`	Assign an IP address to the interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Enter interface mode for xe2
`R2(config-if)#ip address 3.3.3.2/24`	Assign an IP address to the interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode
`R2(config)#router bgp 3`	Enter into BGP router mode
`R2(config-router)#neighbor 2.2.2.2 remote-as 2`	Create static BGP neighbor 2.2.2.2 with remote autonomous system value 2
`R2(config-router)#neighbor 3.3.3.3 remote-as 4`	Create static BGP neighbor 3.3.3.3 with remote autonomous system value 4
`R2(config-router)#address-family ipv4`	Enter address-family mode
`R2(config-router-af)#neighbor 2.2.2.2 activate`	Activate neighbor 2.2.2.2
`R2(config-router-af)#neighbor 3.3.3.3 activate`	Activate neighbor 3.3.3.3
`R2(config-router-af)#redistribute connected`	Advertise the connected network into BGP
`R2(config-router-af)#neighbor 2.2.2.2 route-map r1 in`	Attach the route-map with route discard configured for the neighbor 2.2.2.2 in IN direction
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-router-af)#end`	Exit BGP router mode

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe2`	Enter interface mode for xe2
`R3(config-if)#ip address 3.3.3.3/24`	Assign an IP address to the interface
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration
`R3(config-if)#exit`	Exit interface mode
`R3(config)#interface xe1`	Enter interface mode for xe1
`R3(config-if)#ip address 4.4.4.2/24`	Assign an IP address to the interface
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#router bgp 4`	Enter into BGP router mode
`R3(config-router)#neighbor 3.3.3.2 remote-as 3`	Create static BGP neighbor 3.3.3.2 with remote autonomous system value 3
`R2(config-router)#address-family ipv4`	Enter address-family mode
`R2(config-router-af)#neighbor 3.3.3.2 activate`	Activate neighbor 3.3.3.2
`R3(config-router-af)#redistribute connected`	Advertise the connected network into BGP
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration
`R3(config-router-af)#end`	Exit BGP router.

Validation

```
R2#show running-config bgp
!
router bgp 3
 neighbor 2.2.2.2 remote-as 2
 neighbor 3.3.3.3 remote-as 4
 !
 address-family ipv4 unicast
 redistribute connected
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 route-map r1 in
 neighbor 3.3.3.3 activate
 exit-address-family
!

R2#show ip bgp
BGP table version is 3, local router ID is 2.2.2.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
           l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
       Network          Next Hop          Metric  LocPrf     Weight      Path
*>   1.1.1.0/24         2.2.2.2              0    100          0           2 ?
```

```
*>   2.2.2.0/24         0.0.0.0              0    100        32768           ?
*                       2.2.2.2              0    100          0           2 ?
*>   3.3.3.0/24         0.0.0.0              0    100        32768            ?
*                       3.3.3.3              0    100          0           4 ?
*>   4.4.4.0/24         3.3.3.3              0    100          0           4 ?
Total number of prefixes 4


R2#show running-config prefix-list
!
ip prefix-list p1
 seq 5 permit any
!

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B     1.1.1.0/24 [20/0] is a summary, Null, 00:00:01
C     2.2.2.0/24 is directly connected, xe1, 00:09:57
C     3.3.3.0/24 is directly connected, xe2, 00:09:50
B     4.4.4.0/24 [20/0] via 3.3.3.3, xe2, 00:00:03
C     127.0.0.0/8 is directly connected, lo, 01:18:30

Gateway of last resort is not set

R2#show hsl nh-table
IPv4 FIB 0
0.0.0.0, Null, 00:00:00:00:00:00, Valid  ,
                      1.1.1.0/24, Installed FORWARD
2.2.2.2, xe1, 00:18:23:26:16:45, Valid  , lport:0x8000026, Egress object id:1
00004, refcnt 0, rulecnt 0
3.3.3.3, xe2, 00:18:23:cb:fb:b7, Valid  , lport:0x800002a, Egress object id:1
00003, refcnt 1, rulecnt 0,
                      4.4.4.0/24, Installed FORWARD

IPv4 FIB 1 10.12.29.1, eth0, 00:00:00:00:00:00, Invalid,
                                                , Not Installed TO_CPU

IPv6 FIB 0

IPv6 FIB 1
```

IPv6 Route Discard

Figure 2-2 shows the configuration required to enable static route discard for IPv6.

Configuration

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface lo`	Enter interface mode for loopback
`R1(config-if)#ip address 1.1.1.2/24 secondary`	Assign an IPv4 address to the interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode
`R1(config)#interface xe1`	Enter interface mode for xe1
`R1(config-if)#ipv6 address 2001::2/64`	Assign an IPv6 address to the interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode
`R1(config)#interface xe2`	Enter interface mode for xe2
`R1(config-if)#ipv6 address 1001::2/64`	Assign an IPv6 address to the interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode
`R1(config)#router bgp 2`	Enter BGP router mode
`R1(config-router)#bgp router-id 1.1.1.1`	Specify router ID
`R1(config-router)#neighbor 2001::3 remote-as 3`	Create static BGP neighbor 2001::3 with remote autonomous system value 3
`R1(config-router)#address-family ipv6 unicast`	Enter address family IPv6 unicast mode
`R1(config-router-af)#neighbor 2001::3 activate`	Activate neighbor in IPv6 address family
`R1(config-router-af)#redistribute connected`	Advertise the connected network into BGP
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-router)#end`	Exit BGP router mode

R2

`R2#configure terminal`	Enter configure mode
`R2(config)#interface lo`	Enter interface mode for loopback
`R2(config-if)#ip address 2.2.2.2/24 secondary`	Assign an IPV4 address

`R2(config-router-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode
`R2(config)#ipv6 prefix-list p1 permit any`	Configure IPv6 prefix list.
`R2(config)#route-map r1`	Enter route-map mode.
`R2(config-route-map)#match ipv6 address prefix-list p1`	Configure match IPv6 prefix list p1
`R2(config-route-map)#set interface null0`	Set the interface to null0
`R2(config-router-map)#commit`	Commit the candidate configuration to the running configuration
`R2(config-route-map)#exit`	Exit route-map mode
`R2(config)#interface xe1`	Enter interface mode for xe1
`R2(config-if)#ipv6 address 2001::3/64`	Assign ipv6 address to the interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode
`R2(config)#interface xe2`	Enter interface mode for xe2
`R2(config-if)#ipv6 address 3001::2/64`	Assign an IPv6 address to the interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode
`R2(config)#router bgp 3`	Enter BGP router mode
`R2(config-router)#bgp router-id 2.2.2.2`	Specify router ID
`R2(config-router)#neighbor 2001::2 remote-as 2`	Create static BGP neighbor 2001::2 with remote autonomous system value 2
`R2(config-router)#neighbor 3001::3 remote-as 4`	Create static BGP neighbor 3001::3 with remote autonomous system value 4
`R2(config-router)#address-family ipv6 unicast`	Enter address family IPv6 unicast mode
`R2(config-router-af)#redistribute connected`	Advertise the connected network into BGP
`R2(config-router-af)#neighbor 2001::2 activate`	Activate the neighbor in IPv6 address family
`R2(config-router-af)#neighbor 3001::3 activate`	Activate the neighbor in IPv6 address family
`R2(config-router-af)#neighbor 2001::2 route-map r1 in`	Attach the route-map with route discard configured for the neighbor 2001::2 in IN direction
`R2(config-router-af)#commit`	Commit the candidate configuration to the running configuration
`R2(config-router)#end`	Exit BGP router mode

R3

`R3#configure terminal`	Enter configure mode
`R3(config)#interface lo`	Enter interface mode for loopback
`R3(config-if)#ip address 3.3.3.2/24 secondary`	Assign an IPV4 address to the interface

R3(config-if)#commit	Commit the candidate configuration to the running configuration
R3(config-if)#exit	Exit interface mode
R3(config)#interface xe1	Enter interface mode for xe1
R3(config-if)#ipv6 address 4001::2/64	Assign an IPv6 address to the interface
R3(config-if)#commit	Commit the candidate configuration to the running configuration
R3(config-if)#exit	Exit interface mode
R3(config)#interface xe2	Enter interface mode for xe2
R3(config-if)#ipv6 address 3001::3/64	Assign an IPv6 address to the interface
R3(config-if)#commit	Commit the candidate configuration to the running configuration
R3(config-if)#exit	Exit interface mode
R3(config)#router bgp 4	Enter into BGP router mode
R3(config-router)#bgp router-id 3.3.3.3	Specify router ID
R3(config-router)#neighbor 3001::2 remote-as 3	Create static BGP neighbor 3001::2 with remote autonomous system value 3
R3(config-router)#address-family ipv6 unicast	Enter address family IPv6 unicast mode
R3(config-router-af)#neighbor 3001::2 activate	Activate neighbor in IPv6 address family
R3(config-router-af)#redistribute connected	Advertise the connected network into BGP
R3(config-router-af)#commit	Commit the candidate configuration to the running configuration
R3(config-router)#end	Exit BGP router mode.

Validation

```
R2#show running-config bgp
!
router bgp 3
 bgp router-id 2.2.2.2
 neighbor 2001::2 remote-as 2
 neighbor 3001::3 remote-as 4
 !
 address-family ipv6 unicast
 redistribute connected
 neighbor 2001::2 activate
 neighbor 2001::2 route-map r1 in
 neighbor 3001::3 activate
 exit-address-family
!
R2#show  bgp ipv6
BGP table version is 3, local router ID is 2.2.2.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
  Network    Next Hop                              Metric  LocPrf  Weight  Path
*> 1001::/64 2001::2(fe80::218:23ff:fede:eecf)
                                                     0      100      0        2?
*> 2001::/64  ::                                     0      100    32768       ?
*             2001::2(fe80::218:23ff:fede:eecf)
                                                     0      100      0        2?
*> 3001::/64  ::                                     0      100     32768      ?
*             3001::3(fe80::eef4:bbff:fe84:781b)  0      100       0       4? *> 4001::/
64 3001::3(fe80::eef4:bbff:fe84:781b)  0      100       0       4?

Total number of prefixes 4


R2#show running-config ipv6 prefix-list
!
ipv6 prefix-list p1 seq 5 permit any
!

R2#

R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 00:56:44
B      1001::/64 [20/0] via ::, Null, 00:00:04
C      2001::/64 via ::, xe1, 00:13:03
C      3001::/64 via ::, xe2, 00:12:56
B      4001::/64 [20/0] via fe80::eef4:bbff:fe84:781b, xe50, 00:00:02
C      fe80::/64 via ::, xe50, 00:46:14


R2#show hsl nh-table
IPv4 FIB 0

IPv4 FIB 1
10.12.29.1, eth0, 00:00:00:00:00:00, Invalid,
                     , Not Installed TO_CPU

IPv6 FIB 0
  , Null, 00:00:00:00:00:00, Valid  ,
                     1001::/64, Installed FORWARD
2001::2, xe1, 00:18:23:de:ee:cf, Valid  , lport:0x8000034, Egress object id:100003,
refcnt 0, rulecnt 0
3001::3, xe2, ec:f4:bb:84:78:1b, Valid  , lport:0x8000032, Egress object id:100004,
refcnt 0, rulecnt 0
```

```
fe80::218:23ff:fede:eecf, xe52, 00:18:23:de:ee:cf, Valid  , lport:0x8000034, Egress
object id:100003, refcn
t 0, rulecnt 0
fe80::eef4:bbff:fe84:781b, xe50, ec:f4:bb:84:78:1b, Valid  , lport:0x8000032, Egress
object id:100004, refc
nt 1, rulecnt 0,
                      4001::/64, Installed FORWARD

IPv6 FIB 1
```

CHAPTER 3 Static Route Object Tracking using IP SLA

Overview

Static Route Object Tracking feature tracks the state of an object's reachability using IP SLA. Client processes like Virtual Router Redundancy Protocol (VRRP) or RIB can register their interest in tracking objects and receive notifications when a state change occurs.

IP Service-Level Assurance (SLA) is a protocol used for analyzing IP service levels in applications and services. It employs active traffic-monitoring technology to continuously monitor network traffic. IP SLA utilizes Internet Control Message Protocol (ICMP) pings to detect link failures and notifies registered clients of such events.

The Object Tracking feature ensures clear separation between the objects to be tracked and the actions taken by clients when a tracked object's state changes. Multiple clients, including VRRP and RIB, can register their interest in tracking the same object, each with distinct actions upon state changes. This Tracking feature is available in OAMD.

Each tracked object is assigned a unique number in the tracking command, which client processes use to monitor specific objects. The tracking process monitors events from these tracked objects and records value changes. Any changes in tracked object values are communicated to interested client processes, either immediately or with a specified delay. The object values are reported as either `up` or `down`.

Topology

Figure 3-3 represents a sample network topology featuring RTR1, RTR2, RTR3, and RTR4, used to demonstrate the static route object tracking feature.

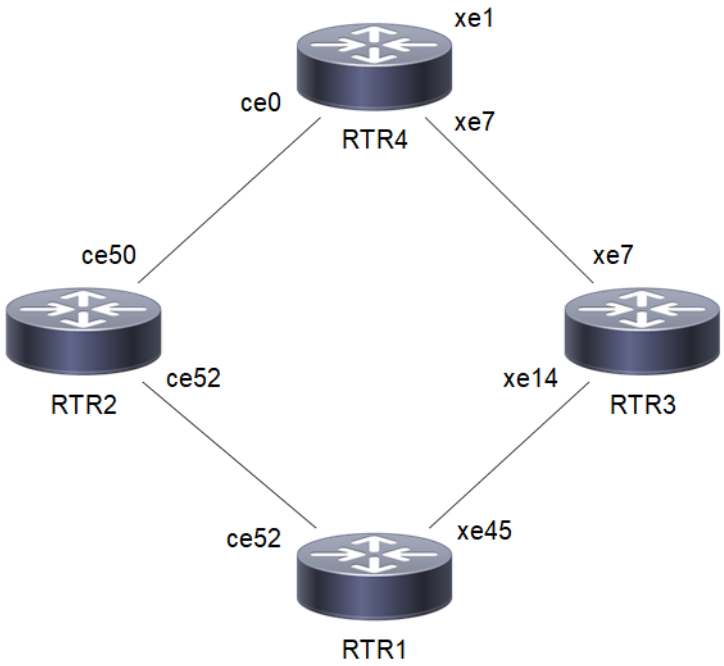

Figure 3-3: Static Route Object Tracking

IPv4 Configuration

The following configuration commands provide an illustrative example of setting up various network parameters on the routers (RTR1, RTR2, RTR3, and RTR4). These commands showcase the configuration of IPv4 addresses, the implementation of IP SLA for network monitoring, the establishment of static routes, and the use of time ranges. Each command serves a specific purpose in configuring and fine-tuning the router's behavior within the network. This example demonstrates a basic network setup with key components for efficient network management and monitoring.

RTR1

`RTR1#configure terminal`	Enter Configure mode.
`RTR1(config)#interface ce52`	Enter interface mode.
`RTR1(config-if)#ip address 2.2.2.1/24`	Configure IPv4 address.
`RTR1(config-if)#exit`	Exit interface mode.
`RTR1(config)#interface xe45`	Enter interface mode.
`RTR1(config-if)#ip address 4.4.4.1/24`	Configure IPv4 address.
`RTR1(config-if)#exit`	Exit interface mode.
`RTR1(config)#ip sla 1`	Configure IP SLA with a unique number.
`RTR1(config-ip-sla)#icmp-echo ipv4 2.2.2.2 source-interface ce52`	Configure the ICMP-echo using destination IPv4 address and source interface name.
`RTR1(config-ip-sla-echo)#threshold 1000`	Configure the threshold value.
`RTR1(config-ip-sla-echo)#timeout 1000`	Configure the timeout value.
`RTR1(config-ip-sla-echo)#frequency 5`	Configure the frequency value.

RTR1(config-ip-sla-echo)#exit	Exit IP SLA echo mode.
RTR1(config-ip-sla)#exit	Exit IP SLA mode.
RTR1(config)#time-range tr1	Configure a time-range.
RTR1(config-tr)#start-time 11:22 3 july 2021	Configure a start-time.
RTR1(config-tr)#end-time after 200	Configure end-time.
RTR1(config-tr)#exit	Exit time-range mode.
RTR1(config)#ip sla schedule 1 time-range tr1	Schedule an IP SLA measurement.
RTR1(config)#track 1 ip sla 1 reachability	Configure the track ID for IP SLA.
RTR1(config-object-track)#exit	Exit object track mode.
RTR1(config)#ip route 3.3.3.0/24 2.2.2.2 track 1	Configure the static route with the nexthop address.
RTR1(config)#ip route 5.5.5.0/24 4.4.4.2	Configure the static route with the nexthop address.
RTR1(config)#ip route 6.6.6.0/24 2.2.2.2 track 1	Configure the static route with the nexthop address and track ID.
RTR1(config)#ip route 6.6.6.0/24 4.4.4.2 10	Configure the static route with the nexthop address and delay.
RTR1(config)#commit	Commit the candidate configuration to the running configuration.
RTR1(config)#exit	Exit configure mode.

RTR2

RTR2#configure terminal	Enter configure mode.
RTR2(config)#interface ce50	Enter interface mode.
RTR2(config-if)#ip address 3.3.3.1/24	Configure IPv4 address.
RTR2(config-if)#exit	Exit interface mode.
RTR2(config)#interface ce52	Enter interface mode.
RTR2(config-if)#ip address 2.2.2.2/24	Configure IPv4 address.
RTR2(config-if)#exit	Exit interface mode.
RTR2(config)#ip route 6.6.6.0/24 3.3.3.2	Configure the static route with the nexthop address.
RTR2(config)#commit	Commit the candidate configuration to the running configuration.
RTR2(config)#exit	Exit configure mode.

RTR3

RTR3#configure terminal	Enter configure mode.
RTR3(config)#interface xe7	Enter interface mode.
RTR3(config-if)#ip address 5.5.5.1/24	Configure IPv4 address.
RTR3(config-if)#exit	Exit interface mode.
RTR3(config)#interface xe14	Enter interface mode.
RTR3(config-if)#ip address 4.4.4.2/24	Configure IPv4 address.

`RTR3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`RTR3(config-if)#exit`	Exit interface mode.
`RTR3(config)#ip route 6.6.6.0/24 5.5.5.2`	Configure the static route with the nexthop address.
`RTR3(config)#commit`	Commit the candidate configuration to the running configuration.
`RTR3(config)#exit`	Exit configure mode.

RTR4

`RTR4#configure terminal`	Enter configure mode.
`RTR4(config)#interface ce0`	Enter interface mode.
`RTR4(config-if)#ip address 3.3.3.2/24`	Configure IPv4 address.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#interface xe1`	Enter interface mode.
`RTR4(config-if)#ip address 6.6.6.6/24`	Configure IPv4 address.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#interface xe7`	Enter interface mode.
`RTR4(config-if)#ip address 5.5.5.2/24`	Configure IPv4 address.
`RTR4(config-if)#commit`	Commit the candidate configuration to the running configuration.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#ip route 2.2.2.0/24 3.3.3.1`	Configure the static route with the nexthop address.
`RTR4(config)#ip route 4.4.4.0/24 5.5.5.1`	Configure the static route with the nexthop address.
`RTR4(config)#commit`	Commit the candidate configuration to the running configuration.
`RTR4(config)#exit`	Exit configure mode.

Validation

RTR1

Below is the routing and tracking information for RTR1, which includes details about its IP SLA reachability tracking and static routes.

```
RTR1#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47

RTR1#show ip route track-table
 ip route 3.3.3.0 255.255.255.0 2.2.2.2 track 1 state is [up]
 ip route 6.6.6.0 255.255.255.0 2.2.2.2 track 1 state is [up]

RTR1#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, xe1, 00:55:39
C            2.2.2.0/24 is directly connected, ce52, 00:55:38
S            3.3.3.0/24 [1/0] via 2.2.2.2, ce52, 00:00:03
C            4.4.4.0/24 is directly connected, xe45, 00:49:50
S            5.5.5.0/24 [1/0] via 4.4.4.2, xe45, 00:08:12
S            6.6.6.0/24 [1/0] via 2.2.2.2, ce52, 00:00:03
C            127.0.0.0/8 is directly connected, lo, 6d23h24m

Gateway of last resort is not set
```

RTR2

```
RTR2#show running-config interface ce52
!
interface ce52
 ip address 2.2.2.2/24
!
```

Now remove the IPv4 address configuration from interface `ce52` on `RTR2` and commit the changes made in the configuration to the running configuration.

```
RTR2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
RTR2(config)#interface ce52
RTR2(config-if)#no ip address
RTR2(config-if)#commit
```

RTR1

The following output from RTR1 routing table and tracking details indicates changes in the reachability of certain routes and related IP SLA measurements:

```
RTR1#sh ip route track-table
2019 Mar 14 14:55:14.350 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_THRESHOLD_CROSSED_2]: IP
SLA 1, response packet 15 exceeds the threshold 10000 ms
2019 Mar 14 14:55:14.351 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_TIMEOUT_2]: IP SLA 1,
response packet 15 exceeds the timeout 10000 ms
2019 Mar 14 14:55:14.351 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_DISCONNECT_2]: Stopping
IP SLA measurement for sla(1) as timeout(10000) observed for packet(15)


 ip route 3.3.3.0 255.255.255.0 2.2.2.2 track 1 state is [down]
 ip route 6.6.6.0 255.255.255.0 2.2.2.2 track 1 state is [down]
```

The IP SLA measurement for `SLA 1` experienced threshold exceedances and timeouts, resulting in changes to the reachability status of the tracked routes. The routes associated with IP SLA tracking, such as `3.3.3.0/24 and 6.6.6.0/24`, are marked as `DOWN`.

```
RTR1#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is DOWN
    5 changes, last change : 2019 Mar 14 14:55:14

2019 Mar 14 14:55:19.352 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_THRESHOLD_CROSSED_2]: IP
SLA 1, response packet 16 exceeds the threshold 10000 ms
2019 Mar 14 14:55:19.353 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_TIMEOUT_2]: IP SLA 1,
response packet 16 exceeds the timeout 10000 ms
2019 Mar 14 14:55:19.353 : R1 : OAM : CRITI : [IPSLA_ICMP_ECHO_DISCONNECT_2]: Stopping
IP SLA measurement for sla(1) as timeout(10000) observed for packet(16)

RTR1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, xe1, 00:56:16
C            2.2.2.0/24 is directly connected, ce52, 00:56:15
C            4.4.4.0/24 is directly connected, xe45, 00:50:27
S            5.5.5.0/24 [1/0] via 4.4.4.2, xe45, 00:08:49
S            6.6.6.0/24 [10/0] via 4.4.4.2, xe45, 00:00:07
C            127.0.0.0/8 is directly connected, lo, 6d23h24m

Gateway of last resort is not set
```

When the IP SLA operation determines that 6.6.6.6 is reachable, the route appears in the routing table, allowing traffic to be routed to that destination. This mechanism provides dynamic routing based on the real-time reachability of the specified IPv4 address.

```
RTR1#ping 6.6.6.6
Press CTRL+C to exit
PING 6.6.6.6 (6.6.6.6) 56(84) bytes of data.
64 bytes from 6.6.6.6: icmp_seq=1 ttl=63 time=0.713 ms
64 bytes from 6.6.6.6: icmp_seq=2 ttl=63 time=0.658 ms
64 bytes from 6.6.6.6: icmp_seq=3 ttl=63 time=0.531 ms
64 bytes from 6.6.6.6: icmp_seq=4 ttl=63 time=0.505 ms

RTR1#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
```

```
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  1.1.1.0/24 is directly connected, xe1, 00:56:58
C    *>  2.2.2.0/24 is directly connected, ce52, 00:56:57
S        3.3.3.0/24 [1/0] via 2.2.2.2, ce52, 00:01:22
C    *>  4.4.4.0/24 is directly connected, xe45, 00:51:09
S    *>  5.5.5.0/24 [1/0] via 4.4.4.2, xe45, 00:09:31
S    *>  6.6.6.0/24 [10/0] via 4.4.4.2, xe45, 00:00:49
S        6.6.6.0/24 [1/0] via 2.2.2.2, ce52, 00:01:22
C    *>  127.0.0.0/8 is directly connected, lo, 6d23h25m


RTR1#show ip sla summary
IP SLA Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                       (usec)        Code        Run
---------------------------------------------------------------------
*1      icmp-echo   2.2.2.2               14000         OK        2019 Mar 14 14:56:26
```

RTR2

Now configure the IPv4 address from interface `ce52` on `RTR2` and commit the changes made in the configuration to the running configuration.

```
RTR2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
RTR2(config)#interface ce52
RTR2(config-if)#ip address 2.2.2.2/24
RTR2(config-if)#commit
```

RTR1

By configuring the IPv4 address on `RTR2` `interface` `ce52` brings up the interface and allows RTR1 to reach RTR2 via the IPv4 address `2.2.2.2/24`, leading to the change in the IP SLA tracking state and the associated static route states to `UP`.

```
RTR1#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    10 changes, last change : 2019 Mar 14 14:56:32

RTR1#show ip route track-table
 ip route 3.3.3.0 255.255.255.0 2.2.2.2 track 1 state is [up]
 ip route 6.6.6.0 255.255.255.0 2.2.2.2 track 1 state is [up]

RTR1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, xe1, 00:57:39
C            2.2.2.0/24 is directly connected, ce52, 00:57:38
S            3.3.3.0/24 [1/0] via 2.2.2.2, ce52, 00:00:11
C            4.4.4.0/24 is directly connected, xe45, 00:51:50
S            5.5.5.0/24 [1/0] via 4.4.4.2, xe45, 00:10:12
S            6.6.6.0/24 [1/0] via 2.2.2.2, ce52, 00:00:11
C            127.0.0.0/8 is directly connected, lo, 6d23h26m

RTR1#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  1.1.1.0/24 is directly connected, xe1, 01:14:49
C    *>  2.2.2.0/24 is directly connected, ce52, 01:14:48
S    *>  3.3.3.0/24 [1/0] via 2.2.2.2, ce52, 00:17:21
C    *>  4.4.4.0/24 is directly connected, xe45, 01:09:00
S    *>  5.5.5.0/24 [1/0] via 4.4.4.2, xe45, 00:27:22
S    *>  6.6.6.0/24 [1/0] via 2.2.2.2, ce52, 00:17:21
S        6.6.6.0/24 [10/0] via 4.4.4.2, xe45, 00:17:25
C    *>  127.0.0.0/8 is directly connected, lo, 6d23h43m
```

IPv6 Configuration

The following configuration commands provide an illustrative example of setting up various network parameters on the routers (RTR1, RTR2, RTR3, and RTR4). These commands showcase the configuration of IPv6 addresses, the implementation of IP SLA for network monitoring, the establishment of static routes, and the use of time ranges. Each command serves a specific purpose in configuring and fine-tuning the router's behavior within the network. This example demonstrates a basic network setup with key components for efficient network management and monitoring.

RTR1

`RTR1#configure terminal`	Enter Configure mode.
`RTR1(config)#interface ce52`	Enter interface mode.
`RTR1(config-if)#ip address 1000::1/64`	Configure IPv6 address.

`RTR1(config-if)#exit`	Exit interface mode.
`RTR1(config)#interface xe45`	Enter interface mode.
`RTR1(config-if)#ip address 4000::2/64`	Configure IPv6 address.
`RTR1(config-if)#exit`	Exit interface mode.
`RTR1(config)#ip sla 1`	Configure IP SLA with a unique number.
`RTR1(config-ip-sla)#icmp-echo ipv6 2000::1 source-interface ce52`	Configure the ICMP-echo using destination IPv6 address and source interface name.
`RTR1(config-ip-sla-echo)#threshold 1000`	Configure the threshold value.
`RTR1(config-ip-sla-echo)#timeout 1000`	Configure the timeout value.
`RTR1(config-ip-sla-echo)#frequency 5`	Configure the frequency value.
`RTR1(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`RTR1(config-ip-sla)#exit`	Exit IP SLA mode.
`RTR1(config)#time-range tr1`	Configure a time-range.
`RTR1(config-tr)#start-time 11:22 3 july 2021`	Configure a start-time.
`RTR1(config-tr)#end-time after 200`	Configure end-time.
`RTR1(config-tr)#exit`	Exit time-range mode.
`RTR1(config)#ip sla schedule 1 time-range tr1`	Schedule an IP SLA measurement.
`RTR1(config)#track 1 ip sla 1 reachability`	Configure the track ID for IP SLA.
`RTR1(config-object-track)#exit`	Configure object track mode.
`RTR1(config)#ipv6 route 3000::0/64 1000::2 track 1`	Configure the static route with the nexthop address.
`RTR1(config)#ipv6 route 3333::1/128 1000::2 track 1`	Configure the static route with the nexthop address and track ID.
`RTR1(config)#ipv6 route 3333::1/128 1000::2 10`	Configure the static route with the nexthop address and delay.
`RTR1(config)#commit`	Commit the candidate configuration to the running configuration
`RTR1(config)#exit`	Exit configure mode

RTR2

`RTR2#configure terminal`	Enter configure mode.
`RTR2(config)#interface ce50`	Enter interface mode.
`RTR2(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address.
`RTR2(config-if)#exit`	Exit interface mode.
`RTR2(config)#interface ce52`	Enter interface mode.
`RTR2(config-if)#ipv6 address 2000::1/64`	Configure IPv6 address.
`RTR2(config-if)#exit`	Exit interface mode.
`RTR2(config)#ipv6 route 3000::0/64 2000::2`	Configure the static route with the nexthop address.

`RTR2(config)#commit`	Commit the candidate configuration to the running configuration.
`RTR2(config)#exit`	Exit configure mode.

RTR3

`RTR3#configure terminal`	Enter configure mode.
`RTR3(config)#interface xe7`	Enter interface mode.
`RTR3(config-if)#ipv6 address 2000::2/64`	Configure IPv6 address.
`RTR3(config-if)#exit`	Exit interface mode.
`RTR3(config)#interface xe14`	Enter interface mode.
`RTR3(config-if)#ipv6 address 3000::1/64`	Configure IPv6 address.
`RTR3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`RTR3(config-if)#exit`	Exit interface mode.
`RTR3(config)#ipv6 route 4000::0/64 3000::2`	Configure the static route with the nexthop address.
`RTR3(config)#ipv6 route 1000::0/64 2000::1`	Configure the static route with the nexthop address.
`RTR3(config)#commit`	Commit the candidate configuration to the running configuration.
`RTR3(config)#exit`	Exit configure mode.

RTR4

`RTR4#configure terminal`	Enter configure mode.
`RTR4(config)#interface ce0`	Enter interface mode.
`RTR4(config-if)#ipv6 address 3000::2/64`	Configure IPv6 address.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#interface xe1`	Enter interface mode.
`RTR4(config-if)#ipv6 address 4000::1/64`	Configure IPv6 address.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#interface xe7`	Enter interface mode.
`RTR4(config-if)#ipv6 address 5000::2/64`	Configure IPv6 address.
`RTR4(config-if)#commit`	Commit the candidate configuration to the running configuration.
`RTR4(config-if)#exit`	Exit interface mode.
`RTR4(config)#ipv6 route 1000::0/64 4000::2`	Configure the static route with the nexthop address.
`RTR4(config)#ipv6 route 2000::0/64 3000::1`	Configure the static route with the nexthop address.
`RTR4(config)#commit`	Commit the candidate configuration to the running configuration.
`RTR4(config)#exit`	Exit configure mode.

Validation

RTR1

Below is the routing and tracking information for RTR1, which includes details about its IP SLA reachability tracking and static routes.

```
RTR1#sh track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47

RTR1#sh ip route track-table
 ipv6 route 3000::0/64 1000::2 track 1 state is [up]
 ipv6 route 3333::1/128 1000::2 track 1 state is [up]

RTR1#sh ip sla summary
IP SLA Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                       (usec)        Code        Run
------------------------------------------------------------------
*1      icmp-echo   2000::2               16000         OK        2019 Mar 11 1
6:11:40
------------------------------------------------------------------

RTR1#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C       1000::/64 via ::, ce52, 00:00:15
C       4000::/64 via ::, ce0, 00:00:15
S       2000::/64 [1/0] via 1000::2, ce52, 00:00:15
S       3333::1/128 [1/0] via 1000::2, ce52, 00:00:15

Gateway of last resort is not set
```

CHAPTER 4 Route Monitor

Overview

Object Tracking provides a mechanism for tracking the reachability status of objects, such as IP status, using Internet Protocol Service Level Agreement (IP SLA). This feature empowers users to monitor the state of these objects and make decisions based on their status. It permits the configuration of multiple track objects on interfaces, delivering flexibility in managing network link status.

Feature Characteristics

Object Tracking establishes a distinct separation between the tracked objects and the actions initiated by a client when there's a change in the state of a tracked object. Users can configure object tracking types as `any` or `all` on the interface, alongside track IDs that specify which statuses to monitor. Modify the interface's link status to either `up` or `down` based on the selected track type and the statuses of the associated track IDs.

When using `Track type all`, the feature performs a `Boolean AND` operation, requiring every object configured on the interface to be in an `up` state for the interface itself to be considered `up`. If any of these objects are not in an `up` state, the interface is set to d`own`.

Conversely, `Track type any` operates as a `Boolean OR` function, necessitating that at least one object configured on the interface must be in an `up` state for the interface to remain `up`. If none of the tracked objects are in an `up` state, the interface is marked as `down`.

Benefits

Users can ensure network reliability by defining specific tracking criteria and actions, allowing them to take appropriate measures when tracked objects experience status change. This contributes to improved network management and performance.

Prerequisites

Before configuring and utilizing Object Tracking, ensure the following prerequisites:

Track IDs: Users must define and configure the track IDs and corresponding objects they want to track for reachability. These track IDs are essential for the feature to work effectively. Deleting all track IDs from the interface will bring the interface up if it was previously down.

Interface Configuration: The feature involves configuring track types on interfaces. Therefore, ensuring that the interfaces are correctly configured and operational is important. In cases where an interface has both object tracking configurations and next-hop reachability, deleting the object tracking configurations is necessary to bring the interface back up if it goes down.

Object Tracking Criteria: Define the specific criteria and conditions for tracking an object’s reachability, such as IP status, using IP SLA.

© 2025 IP Infusion Inc. Proprietary

Configuration

The below topology illustrates a network configuration involving three routers, R1, R2, and R3, with a central device referred to as the Device Under Test (DUT) positioned in the middle. This topology represents a linear or sequential network structure that showcases the Route Monitor feature.

Topology

A series of configurations were implemented on routers R1, R2, and R3, as well as on the DUT, to showcase the functionality of the Route Monitor feature. The objective was to demonstrate the configuration of network routers to monitor the reachability status of specific IPv4 and IPv6 addresses using IP SLA and illustrate that these configurations can work in conjunction with the Route Monitor feature to enable informed decisions based on the reachability status of tracked objects.

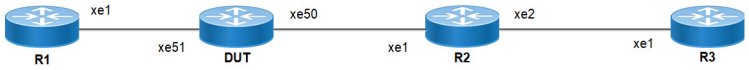

Figure 4-4: Route Monitor Topology

IPv4 Configuration

DUT

Use the following configuration to set up an IP SLA and enable object tracking on a network device. These commands assign IPv4 addresses to interfaces, configure specific IP SLA parameters such as threshold, timeout, and frequency, create a time-range for scheduling measurements, and establish static routes with nexthop addresses. Configure object tracking to monitor the reachability of tracked objects. These configurations highlight the versatility and functionality of the network device by allowing it to monitor IPv4 addresses, make decisions based on object tracking, and optimize network operations.

`DUT#configure terminal`	Enter configure mode.
`DUT(config)#interface xe50`	Enter interface mode `xe50`.
`DUT(config-if)#ip address 2.2.2.1/24`	Assign the IP address `2.2.2.1` with a subnet mask of `/24` to interface `xe50`.
`DUT(config-if)#exit`	Exit interface mode `xe50`.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#ip address 1.1.1.2/24`	Assign the IP address 1.1.1.2 with a subnet mask of /24 to interface xe51.
`DUT(config-if)#exit`	Exit interface mode `xe51`.
`DUT(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`DUT(config-ip-sla)#icmp-echo ipv4 3.3.3.1 source-interface xe50`	Configure the SLA to send ICMP echo requests to destination IPv4 address `3.3.3.1` using interface `xe50` as the source.
`DUT(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send ICMP echo packets every 5 seconds.
`DUT(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.

© 2025 IP Infusion Inc. Proprietary

`DUT(config-ip-sla)#exit`	Exit IP SLA mode.
`DUT(config)#time-range tr1`	Create a time range named `tr1`.
`DUT(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on July 3, 2021`.
`DUT(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`DUT(config-tr)#exit`	Exit time-range mode.
`DUT(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation 1 to run within the specified time range `tr1`.
`DUT(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`DUT(config-object-track)#exit`	Exit object track mode.
`DUT(config)#ip route 3.3.3.0/24 2.2.2.2 track 1`	Add a static route for the destination network `3.3.3.0/24` with next-hop IP `2.2.2.2`, tracked by tracking object `1`.
`DUT(config)#ip route 5.5.5.0/24 1.1.1.2`	Add a static route for the destination network `5.5.5.0/24` with next-hop IP `1.1.1.2`.
`DUT(config)#ip route 6.6.6.0/24 2.2.2.2 track 1`	Add a static route for the destination network `6.6.6.0/24` with next-hop IP `2.2.2.2`, tracked by tracking object `1`.
`DUT(config)#ip route 6.6.6.0/24 1.1.1.2 10`	Add a static route for the destination network `6.6.6.0/24` with next-hop IP `1.1.1.2` and a metric of `10`.
`DUT(config)#commit`	Commit the candidate configuration to the running configuration.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`DUT(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`DUT(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`DUT(config-if)#exit`	Exit interface mode.
`DUT(config)#exit`	Exit configure mode.

By configuring the routes below, R1, R2, and R3 effectively forward network traffic to its designated destinations within the network. These configurations actively contribute to efficient routing operations and ensure network traffic reaches its targets.

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#ip address 1.1.1.1/24`	Assign the IP address `1.1.1.1` with a subnet mask of `/24` to interface `xe1`.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit interface mode `xe1`.
`R1(config)#ip route 2.2.2.0/24 1.1.1.2`	Add a static route for the destination network `2.2.2.0/24` with next-hop IP `1.1.1.2`.
`R1(config)#ip route 3.3.3.0/24 1.1.1.2`	Add a static route for the destination network `3.3.3.0/24` with next-hop IP `1.1.1.2`.

`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit configure mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe1`	Enter interface mode `xe1`.
`R2(config-if)#ip address 2.2.2.2/24`	Assign the IP address `2.2.2.2` with a subnet mask of `/24` to interface `xe1`.
`R2(config-if)#exit`	Exit interface mode `xe1`.
`R2(config)#interface xe2`	Enter interface mode `xe2`.
`R2(config-if)#ip address 3.3.3.1/24`	Assign the IP address `3.3.3.1` with a subnet mask of `/24` to interface `xe2`.
`R2(config-if)#exit`	Exit interface mode `xe2`.
`R2(config)#ip route 1.1.1.0/24 2.2.2.1`	Add a static route for the destination network `1.1.1.0/24` with next-hop IP `2.2.2.1`.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit configure mode.

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#ip address 3.3.3.2/24`	Assign the IP address `3.3.3.2` with a subnet mask of `/24` to interface `xe1`.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit interface mode `xe1`.
`R3(config)#ip route 1.1.1.0/24 3.3.3.1`	Add a static route for the destination network `1.1.1.0/24` with next-hop IP `3.3.3.1`.
`R3(config)#ip route 2.2.2.0/24 3.3.3.1`	Add a static route for the destination network `2.2.2.0/24` with next-hop IP `3.3.3.1`.
`R3(config)#commit`	Commit the candidate configuration to the running configuration.
`R3(config)#exit`	Exit configure mode.

Validation

The following show output displays information about the IPv4 route table, IP SLA reachability tracking, and interface status on a network device running OcNOS.

DUT

```
DUT#show track
TRACK Id: 1
  IP SLA 1 reachability
```

```
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47
Track interface : xe51

DUT#show ip route track-table
 ip route 3.3.3.0 255.255.255.0 2.2.2.2 track 1 state is [up]
 ip route 6.6.6.0 255.255.255.0 2.2.2.2 track 1 state is [up]

DUT#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, xe51, 00:55:38
C            2.2.2.0/24 is directly connected, xe50, 00:49:50
S            3.3.3.0/24 [1/0] via 2.2.2.2, xe50, 00:00:03
S            5.5.5.0/24 [1/0] via 1.1.1.2, xe51, 00:08:12
S            6.6.6.0/24 [1/0] via 2.2.2.2, xe50, 00:00:03

Gateway of last resort is not set

DUT#show interface brief xe51

Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
       FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
       CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
       ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
       PD(Min L/B) - Protocol Down Min-Links/Bandwidth
       OTD  - Object Tracking Down
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       Ctl - Control Port (Br-Breakout/Bu-Bundle)


--------------------------------------------------------------------------------
Ethernet  Type    PVID   Mode     Status  Reason  Speed Port ch#  Ctl Br/Bu  Loopbk Interface
--------------------------------------------------------------------------------
xe51       ETH    --    routed   down    OTD     10g   --      No       No
```

IPv6 Configuration

DUT

Use the following configuration to set up an IP SLA and enable object tracking on a network device. These commands assign IPv6 addresses to interfaces, configure specific IP SLA parameters such as threshold, timeout, and frequency, create a time-range for scheduling measurements, and establish static routes with nexthop addresses. Configure object tracking to monitor the reachability of tracked objects. These configurations highlight the versatility and functionality of the network device by allowing it to monitor IPv6 addresses, make decisions based on object tracking, and optimize network operations.

`DUT#configure terminal`	Enter configure mode.
`DUT(config)#interface xe50`	Enter interface mode `xe50`.
`DUT(config-if)#ipv6 address 2000::1/64`	Assign an IPv6 address (`2000::1/64`) to interface `xe50`.
`DUT(config-if)#exit`	Exit interface mode `xe50`.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#ipv6 address 1000::2/64`	Assign an IPv6 address (`1000::2/64`) to interface `xe51`.

`DUT(config-if)#exit`	Exit interface mode `xe51`.
`DUT(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`DUT(config-ip-sla)#icmp-echo ipv6 3000::1 source-interface xe50`	Configure the SLA to send IPv6 ICMP echo requests to destination IPv6 address `3000::1` using interface `xe50` as the source.
`DUT(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send IPv6 ICMP echo packets every 5 seconds.
`DUT(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`DUT(config-ip-sla)#exit`	Exit IP SLA mode.
`DUT(config)#time-range tr1`	Create a time range named `tr1`.
`DUT(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on July 3, 2021`.
`DUT(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`DUT(config-tr)#exit`	Exit time-range mode.
`DUT(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation `1` to run within the specified time range `tr1`.
`DUT(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`DUT(config-object-track)#exit`	Exit object track mode.
`DUT(config)#ipv6 route 3000::0/64 2000::2 track 1`	Add an IPv6 static route for the destination network `3000::0/64` with a next-hop IPv6 `2000::2`, tracked by tracking object `1`.
`DUT(config)#ipv6 route 3333::1/128 1000::1`	Add an IPv6 static route for the destination network `3333::1/128` with next-hop IPv6 `1000::1`.
`DUT(config)#ipv6 route 3333::1/128 2000::2 track 1`	Add an IPv6 static route for the destination network `6.6.6.0/24` with next-hop IPv6 `2000::2`, tracked by tracking object `1`.
`DUT(config)#ipv6 route 3333::1/128 1000::1 10`	Add an IPv6 static route for the destination network `3333::1/128` with next-hop IP `1000::1` and a metric of `10`.
`DUT(config)#commit`	Commit the candidate configuration to the running configuration.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`DUT(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`DUT(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`DUT(config-if)#exit`	Exit interface mode.
`DUT(config)#exit`	Exit configure mode.

By configuring the routes below, R1, R2, and R3 effectively forward network traffic to its designated destinations within the network. These configurations actively contribute to efficient routing operations and ensure network traffic reaches its targets.

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#ipv6 address 1000::1/64`	Assign the IPv6 address `1000::1` with a subnet mask of `/64` to interface `xe1`.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit interface mode `xe1`.
`R1(config)#ipv6 route 2000::0/64 1000::2`	Add an IPv6 static route for the destination network `2000::0/64` with next-hop IPv6 `1000::2`.
`R1(config)#ipv6 route 3000::0/64 1000::2`	Add an IPv6 static route for the destination network `3000::0/64` with next-hop IPv6 `1000::2`.
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit configure mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe1`	Enter interface mode `xe1`.
`R2(config-if)#ipv6 address 2000::2/64`	Assign the IPv6 address `2000::2` with a subnet mask of `/64` to interface `xe1`.
`R2(config-if)#exit`	Exit interface mode `xe1`.
`R2(config)#interface xe2`	Enter interface mode `xe2`.
`R2(config-if)#ipv6 address 3000::1/64`	Assign the IPv6 address `3000::1` with a subnet mask of `/64` to interface `xe2`.
`R2(config-if)#exit`	Exit interface mode `xe2`.
`R2(config)#ipv6 route 1000::0/64 2000::1`	Add an IPv6 static route for the destination network `1000::0/64` with next-hop IPv6 `2000::1`.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit configure mode.

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#ipv6 address 3000::2/64`	Assign the IPv6 address `3000::2` with a subnet mask of `/64` to interface `xe1`.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit interface mode `xe1`.
`R3(config)#ipv6 route 1000::0/64 3000::1`	Add an IPv6 static route for the destination network `1000::0/64` with next-hop IPv6 `3000::1`.

`R3(config)#ipv6 route 2000::0/64 3000::1`	Add an IPv6 static route for the destination network `2000::0/64` with next-hop IPv6 `3000::1`.
`R3(config)#commit`	Commit the candidate configuration to the running configuration.
`R3(config)#exit`	Exit configure mode.

Validation

The following show output displays the information about IP SLA reachability tracking, IPv6 route tables, and interface status on a network device running OcNOS.

DUT

```
DUT#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47
Track interface : xe51

DUT#show ip route track-table
 ipv6 route 3000::0/64 2000::2 track 1 state is [up]
 ipv6 route 3333::1/128 2000::2 track 1 state is [up]

DUT#show ip sla summary
IP SLA Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                       (usec)        Code        Run
------------------------------------------------------------------------
*1      icmp-echo   3000::1                16000         OK        2019 Mar 11 1
6:11:40
------------------------------------------------------------------------

DUT#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:04:46
C       1000::/64 via ::, xe51, 00:02:48
C       2000::/64 via ::, xe50, 00:02:48
S       3000::/64 [1/0] via 2000::2, xe50, 00:02:48
S       3333::1/128 [1/0] via 2000::2, xe50, 00:02:48

DUT#show interface brief xe51

Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
       FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
       CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
       ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
       PD(Min L/B) - Protocol Down Min-Links/Bandwidth
       OTD  - Object Tracking Down
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       Ctl - Control Port (Br-Breakout/Bu-Bundle)


--------------------------------------------------------------------------------
Ethernet   Type     PVID  Mode      Status  Reason  Speed Port Ch #  Ctl Br/Bu  Loopbk
Interface
```

```
-------------------------------------------------------------------------------
xe51        ETH     --    routed     down    OTD     10g   --         No         No
```

Implementation Examples

Here is a practical scenario and use case for Object Tracking implementation:

Link Redundancy: Object Tracking can be used to monitor the reachability of primary and backup network links. If the primary link fails or becomes congested, the system can automatically switch traffic to the backup link, ensuring uninterrupted network connectivity.

Load Balancing: Object Tracking helps optimize load balancing by continuously assessing the health and availability of servers or paths. If a server becomes overloaded or fails, traffic can be intelligently redirected to healthy servers, improving resource utilization and user experience.

Failover Testing and Verification: Object Tracking provides a mechanism for simulating network failures and verifying failover mechanisms. By configuring tracked objects to mimic real-world conditions, network administrators can assess the resilience of their network configurations and ensure they perform as expected during failures.

New CLI Commands

The Route Monitor feature introduces the following configuration commands. For more information, refer to the *Interface Commands, IP Service Level Agreements Commands,* and *Object Tracking Commands* chapters in the System Management Guide.

object-tracking

Use this command to configure track IDs and options on the interfaces.

Use the no parameter with this command to remove the configurations.

These commands configure object tracking on interfaces, with specific track IDs and tracked objects set to determine what gets tracked and affects the interface's status.

The `object-tracking` command provides flexibility, enabling both `all` and `any` tracking behaviors for influencing the interface's status. A maximum of 8 track IDs can be configured per interface. It is possible to configure the same track IDs or options on multiple interfaces.

Command Syntax

```
object-tracking <1-500>
object-tracking <all | any>
no object-tracking <1-500>
no object-tracking <all | any>
```

Parameters

`<1-500>`	Object tracking ID
`all`	Boolean AND operation. Each object configured on the interface must be in an up state for the interface itself to be in an up state; otherwise, it will be brought down.
`any`	Boolean OR operation. At least one object configured on the interface must be in an up state; otherwise, the interface will be brought down.

Default

None

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Example

Here are some example commands for configuring object tracking in the interface mode.

```
OcNOS(config)#int xe5
OcNOS(config-if)#object-tracking 10
OcNOS(config-if)#object-tracking all
OcNOS(config-if)#commit

OcNOS(config-if)#no object-tracking 10
OcNOS(config-if)#no object-tracking all
OcNOS(config-if)#commit
OcNOS(config-if)#exit
```

Troubleshooting

Interface Status: Verify the status of the interface linked with object tracking. If the configured track type is `all`, confirm that all tracked objects are in an `up` state to consider the interface as `up`. In the case of the track type being `any`, ensure that at least one tracked object is `up` to maintain the interface in an `up` state.

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
NSM	Network and Service Management
IP SLA	Internet Protocol Service Level Agreement
DUT	Device Under Test

Glossary

The following provides definitions for key terms used throughout this document.

Object Tracking	A feature that monitors the reachability status of objects, such as IP status, using IP SLA and allows users to take actions based on their status.
Track Object	An object configured for tracking within the Object Tracking feature. These objects can represent specific network components or conditions, such as IP addresses or link statuses.
Track ID	A unique identifier associated with a track object that enables the system to monitor and assess the status of that object.
Track Type	The configuration specifies how the interface's link status should be determined based on the statuses of associated track objects. It can be set to all or any.
Track Type “All”	A track type that uses a Boolean AND function, requiring that all tracked objects be in an up state for the interface to be considered up.
Track Type “Any”	A track type that uses a Boolean OR function, ensuring that at least one tracked object is in an up state for the interface to remain up.

CHAPTER 5 Route-map Continue Configuration

This section contains Route-map continue configuration with BGP.

Overview

The continue clauses allow you to configure and organize more modular policy definitions to reduce the number of policy configurations that are repeated within the same route map.

Continue clause under a route-map provides the capability to execute additional entries in a route map after an entry is executed with a successful match and set clauses. The continue command allows multiple entries to be evaluated within a single route-map. Continue commands can be assigned optional sequence numbers that indicates the order in which clauses are to be evaluated.

Using Continue with Match Clauses

When a `match` clause exists in a route-map with continue clause then `continue` clause is executed only when a successful match occurs. If a `match` clause does not exist in the route-map and if a `continue` clause exists, the `continue` clause will be evaluated and will go to the specified route-map entry. When a successful match occurs and we have a `continue` clause, the route-map executes the `set` clauses and then goes to the specified route-map entry. If a `continue` clause does not exist in the next route map, then the route-map will behave normally. If a `continue` clause exists in the next route-map but a match is not successful, the route-map will not continue and will “fall through” to the next sequence number if one exists

Using Continue with Set Actions

`Set` clauses are executed after the route-map evaluation is done. The `set` clauses are evaluated and executed in the order in which they were configured. `Set` clauses are only executed after a successful match occurs. The `continue` statement proceeds to the specified route-map entry only after configured `set` actions are performed. If a `set` action is configured in the first route-map and then the same `set` action occurs again but with a different value, in a subsequent route-map entry, then the last `set` action will override the previous `set` actions which were configured with the same set command.

Topology

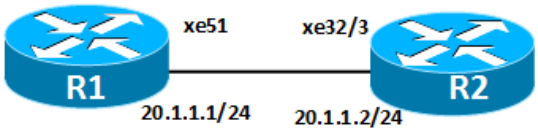

Figure 5-5: Route-map continue

Configure Route-map *continue* on R1

In the below example we will apply route-map `continue` on R1 under BGP 100, with redistributed connected routes from R2 on R1. Here, 10.1.0.0/16 is a superset, while 10.1.1.0/24, 10.1.2.0 /24, 10.1.3.0/24, 10.1.4.0/24, and 10.1.5.0/24 are subsets and will filter PF1, which is a superset and sets several prefixes.

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe51`	Enter interface mode.
`R1(config-if)#ip address 20.1.1.1/24`	Configure the IP address of the interface.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#ip prefix-list PF1`	Configure super set prefix PF1
`R1(config-ip-prefix-list)#seq 5 permit 10.1.0.0/16 le 32`	Permit 10.1.0.0/16 length as 32
`R1(config)#ip prefix-list P1`	Configure subset prefix P1
`R1(config-ip-prefix-list)#seq 5 permit 10.1.1.0/24 le 32`	Permit 10.1.1.0/24 length as 32
`R1(config)#ip prefix-list P2`	Configure subset prefix P2
`R1(config-ip-prefix-list)#seq 5 permit 10.1.2.0/24 le 32`	Permit 10.1.2.0/24 length as 32
`R1(config-ip-prefix-list)#ip prefix-list P3`	Configure subset prefix P3
`R1(config-ip-prefix-list)#seq 5 permit 10.1.3.0/24 le 32`	Permit 10.1.3.0/24 length as 32
`R1(config-ip-prefix-list)#exit`	Exit from prefix-list mode
`R1(config)#route-map myid1 permit 1`	Configure route-map myid1 with sequence number 1
`R1(config-route-map)match ip address prefix-list PF1`	Match for prefix PF1
`R1(config-route-map)continue`	Configure continue command without sequence number
`R1(config-route-map)set metric 10`	Set metric as 10
`R1(config-route-map)set weight 3465789`	Set weight as 3465789
`R1(config-route-map)route-map myid1 permit 2`	Configure route-map myid1 with sequence number 2
`R1(config-route-map)match ip address prefix-list P1`	Match for IP prefix-list P1
`R1(config-route-map)continue 3`	Configure continue with sequence number 3
`R1(config-route-map)set metric 20`	Set metric as 20.
`R1(config-route-map)set origin igp`	Set origin as IGP protocol
`R1(config-route-map)route-map myid1 permit 3`	Configure route-map myid1 with sequence number 3
`R1(config-route-map)match ip address prefix-list P2`	Match for IP prefix-list P2
`R1(config-route-map)continue 4`	Configure continue with sequence number 4
`R1(config-route-map)set metric 30`	Set metric as 30.
`R1(config-route-map)set as-path prepend 600`	Set as-path prepend as 600
`R1(config-route-map)route-map myid1 permit 4`	Configure route-map myid1 with sequence number 4
`R1(config-route-map)match ip address prefix-list P3`	Match for IP prefix-list P3.
`R1(config-route-map)set local-preference 400`	Set local preference as 400
`R1(config-route-map)set weight 400`	Set weight as 400
`R1(config-route-map)end`	End the route-map
`R1(config)#router bgp 100`	Configure bgp process 100
`R1(config-router)#bgp router-id 1.1.1.1`	Configure bgp router id

© 2025 IP Infusion Inc. Proprietary

`R1(config-router)#neighbor 20.1.1.2 remote-as 100`	Configure bgp remote-as 100 with neighbor IP
`R1(config-router)#address-family ipv4`	Enter address-family mode
`R1(config-router-af)#neighbor 20.1.1.2 activate`	Activate neighbor 20.1.1.2
`R1(config-router-af)#neighbor 20.1.1.2 route-map myid1 in`	Configure bgp route-map myid1 as In bound policy with neighbor ip
`R1(config-router-af)#end`	Exit the router bgp mode

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe32/1`	Enter interface mode.
`R2(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface on an interface which is up and running
`R2(config-if)#interface xe32/2`	Enter interface mode.
`R2(config-if)#ip address 10.1.2.1/24`	Configure the IP address of the interface on an interface which is up and running
`R2(config-if)#interface xe32/4`	Enter interface mode.
`R2(config-if)#ip address 10.1.3.1/24`	Configure the IP address of the interface on an interface which is up and running
`R2(config-if)#interface xe17/1`	Enter interface mode.
`R2(config-if)#ip address 10.1.4.1/24`	Configure the IP address of the interface on an interface which is up and running
`R2(config-if)#interface xe21/1`	Enter interface mode.
`R2(config-if)#ip address 10.1.5.1/24`	Configure the IP address of the interface on an interface which is up and running
`R2(config-if)#interface xe32/3`	Enter interface mode.
`R2(config-if)#ip address 20.1.1.2/24`	Configure the IP address on the connected interface.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#router bgp 100`	Configure BGP process 100
`R2(config-router)#bgp router-id 2.2.2.2`	Configure BGP router id
`R2(config-router)#neighbor 20.1.1.1 remote-as 100`	Configure BGP remote-as 100 with neighbor IP
`R2(config-router)#address-family ipv4`	Enter address-familymode
`R2(config-router-af)#neighbor 20.1.1.1 activate`	Activate neighbor20.1.1.1
`R2(config-router-af)#redistribute connected`	Redistribute the connected routes which are 10 networks here.
`R2(config-router-as)#end`	Exit router BGP mode

© 2025 IP Infusion Inc. Proprietary

Validation

R1

```
R1#show ip bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 5
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
20.1.1.2                 4   100  145        177       5      0      0
00:40:05              5

Total number of neighbors 1

Total number of Established sessions 1

Note: Check the prefixes learnt here are 5.

R1#

R1#show ip bgp
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  10.1.1.0/24      20.1.1.2              20        100           3465789  i
*>i  10.1.2.0/24      20.1.1.2              30        100           3465789 600
?
*>i  10.1.3.0/24      20.1.1.2              10        400           400      ?
*>i  10.1.4.0/24      20.1.1.2              10        100           3465789  ?
*>i  10.1.5.0/24      20.1.1.2              10        100           3465789  ?

Total number of prefixes 5
R1#
```

Note: In the above example, 10.1.4.0/24 and 10.1.5.0/24 prefixes will match only on PF1 which is a super set prefix and metric is set as 10, while the 10.1.1.0/24, 10.1.2.0/24 and 10.1.3.0/24 prefixes will match in P1, P2 and P3 prefix-lists and execute the set clauses respectively.

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default
```

```
IP Route Table for VRF "default"
B     10.1.1.0/24 [200/20] via 20.1.1.2, xe51, 00:45:05
B     10.1.2.0/24 [200/30] via 20.1.1.2, xe51, 00:45:05
B     10.1.3.0/24 [200/10] via 20.1.1.2, xe51, 00:45:05
B     10.1.4.0/24 [200/10] via 20.1.1.2, xe51, 00:25:05
B     10.1.5.0/24 [200/10] via 20.1.1.2, xe51, 00:24:35
C     20.1.1.0/24 is directly connected, xe51, 01:00:40
C     127.0.0.0/8 is directly connected, lo, 02:26:41

Gateway of last resort is not set
R1#

R1#show ip bgp route-map myid1
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  10.1.1.0/24      20.1.1.2              20        100           3465789  i
*>i  10.1.2.0/24      20.1.1.2              30        100           3465789 600
?
*>i  10.1.3.0/24      20.1.1.2              10        400           400      ?
*>i  10.1.4.0/24      20.1.1.2              10        100           3465789  ?
*>i  10.1.5.0/24      20.1.1.2              10        100           3465789  ?

Total number of prefixes 5
R1#
R1#
```

R2

```
R2#show ip bgp summary
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
20.1.1.1                 4   100  133        133       3      0      0
00:39:57              0

Total number of neighbors 1

Total number of Established sessions 1
R2#
R2#


R2#show ip bgp
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric    LocPrf       Weight Path
*>   10.1.1.0/24      0.0.0.0              0         100            32768    ?
*>   10.1.2.0/24      0.0.0.0              0         100            32768    ?
*>   10.1.3.0/24      0.0.0.0              0         100            32768    ?
*>   10.1.4.0/24      0.0.0.0              0         100            32768    ?
*>   10.1.5.0/24      0.0.0.0              0         100            32768    ?
*>   20.1.1.0/24      0.0.0.0              0         100            32768    ?

Total number of prefixes 6
R2#

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     10.1.1.0/24 is directly connected, xe32/1, 01:02:22
C     10.1.2.0/24 is directly connected, xe32/2, 01:01:46
C     10.1.3.0/24 is directly connected, xe32/4, 01:02:21
C     10.1.4.0/24 is directly connected, xe17/1, 00:26:52
C     10.1.5.0/24 is directly connected, xe21/1, 00:26:32
C     20.1.1.0/24 is directly connected, xe32/3, 01:02:22
C     127.0.0.0/8 is directly connected, lo, 01:34:40

Gateway of last resort is not set
R2#
```

Route-map Show Commands

```
R1#show running-config route-map
    !
    route-map myid1 permit 1
     match ip address prefix-list PF1
     continue
     set metric 10
     set weight 3465789
    !
    route-map myid1 permit 2
     match ip address prefix-list P1
     continue 3
     set metric 20
     set origin igp
    !
    route-map myid1 permit 3
     match ip address prefix-list P2
     continue 4
     set metric 30
     set as-path prepend 600
```

```
    !
    route-map myid1 permit 4
     match ip address prefix-list P3
     set local-preference 400
     set weight 400
    !
    R1#
R1#show route-map
    route-map myid1, permit, sequence 1
      Match clauses:
        ip address prefix-list: PF1
      Continue clause:  next sequence
      Set clauses:
        metric 10
        weight 3465789
    route-map myid1, permit, sequence 2
      Match clauses:
        ip address prefix-list: P1
      Continue clause:  sequence 3:
      Set clauses:
        metric 20
        origin igp
    route-map myid1, permit, sequence 3
      Match clauses:
        ip address prefix-list: P2
      Continue clause:  sequence 4:
      Set clauses:
        metric 30
        as-path prepend 600
    route-map myid1, permit, sequence 4
      Match clauses:
        ip address prefix-list: P3
      Set clauses:
        local-preference 400
        weight 400
    R1#
```

CHAPTER 6 Unidirectional Link Detection Configuration

This chapter shows a complete configuration to enable UDLD in a simple network topology.

Overview

The purpose of Unidirectional Link Detection protocol (UDLD) is to monitor the physical links and detect when a unidirectional link exists. Upon detection user can either block the port or notify the link status based on the network administrator's configuration.

UDLD works in two different modes:

- Normal mode
- Aggressive mode

Topology

Figure 6-6 shows the topology of the UDLD configuration.

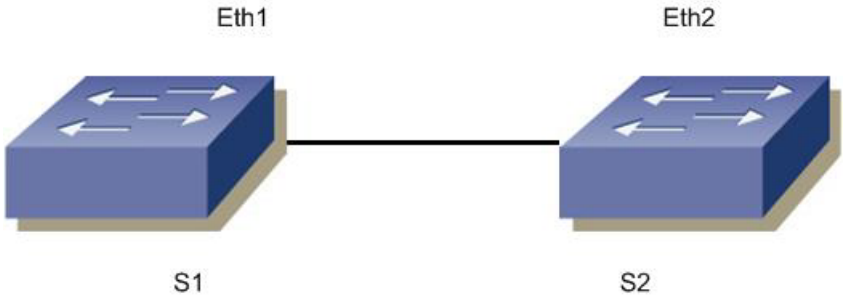

Figure 6-6: UDLD Configuration

S1

`#configure terminal`	Enter configure mode
`(config)#udld enable`	Enable UDLD globally
`(config)#udld message-time 7`	Configure message time for UDLD packets
`(config)#interface eth1`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switch port
`(config-if)#udld state enable`	Enable UDLD on the interface
`(config-if)#udld mode normal`	Configure udld mode as normal or aggressive
`(config-if)#commit`	Commit config.
`(config-if)#exit`	Exit from the interface mode

S2

`#configure terminal`	Enter configure mode.
`(config)#udld enable`	Enable UDLD globally.
`(config)#udld message-time 7`	Configure message time for UDLD packets

© 2025 IP Infusion Inc. Proprietary

(config)#interface eth2	Enter interface mode
(config-if)#switchport	Configure the interface as switch port.
(config-if)#udld state enable	Enable UDLD on the interface.
(config-if)#udld mode normal	Configure udld mode as normal or aggressive
(config-if)#commit	Commit config.
(config-if)#exit	Exit from the interface mode

Validation

```
#show udld
UDLD: Enable
Message Interval(sec) : 7

Port   UDLD Status      Mode      Link-Status
-------------------------------------------------
Eth1   Enable           Normal    Bi-directional
Eth2   Disable          Normal    Unknown
Eth3   Disable          Normal    Unknown
Eth4   Disable          Normal    Unknown
Eth5   Disable          Normal    Unknown
Eth6   Disable          Normal    Unknown
```

Once the links is made Uni-directional, the output of the command Show udld is as follows:

```
#show udld
UDLD: Enable
Message Interval(sec) : 7

Port   UDLD Status      Mode       Link-Status
-------------------------------------------------
Eth1   Enable           Normal     Unidirectional
Eth2   Disable          Normal     Unknown
Eth3   Disable          Normal     Unknown
Eth4   Disable          Normal     Unknown
Eth5   Disable          Normal     Unknown
Eth6   Disable          Normal     Unknown
```

```
#sh running-config
udld Enable
udld message-time 7

#sh running-config in eth1
!
interface eth1
 switchport
 udld state Enable
```

```
!

#sh udld interface eth1
UDLD Status           : Enable
UDLD Mode             : Normal
Link-State            : Unknown
```

For aggressive mode, udld output is as follows:

```
#show udld
    UDLD  : Enable
    Message Interval(sec) : 7


Port   UDLD Status     Mode        Link-Status
-----------------------------------------------
eth1   Enable          Aggressive  Bi-Directional


#sh running config
udld Enable
udld message-time 7

#sh running-config in eth1
  interface eth1
switchport
udld mode Aggressive
udld state Enable
```

CHAPTER 7 Internet Protocol SLA Configuration

Internet Protocol Service Level Agreement (IP SLA) is an active method of monitoring and reliably reporting on network performance. By "active," I refer to the fact that IP SLA will generate and actively monitor traffic continuously across the network. An IP SLA Router is capable of generating traffic and reporting on it in real time

IP SLA can be configured in two parts. There is the IP SLA router, which generates the traffic, and the IP SLA Responder (which can be any device, not just a router). The IP SLA Responder is not required for IP SLA to function, but it does allow for more detailed information gathering and reporting.

After an IP SLAs operation has been configured, you must schedule the operation to begin capturing statistics and collecting error information. When scheduling an operation, it can start immediately or start at a certain month, day, and hour. There is a pending option to set the operation to start at a later time. The pending option is also an internal state of the operation visible through SNMP. The pending state is also used when an operation is a reaction (threshold) operation waiting to be triggered. You can schedule a single ip slas operation or a group of operations at one time.

Note: IP SLA sessions are scaled to 500 sessions on Edgecore AS7316-26XB switches. This limit may vary on other devices based on the device capacity and performance.

Topology

Figure 7-7: IP SLA Topology

Configuration

Configure IP Address

Configure the IP addresses on the PE-1, P routers.

PE-1

`#configure terminal`	Enter configure mode.
`(config)#interface xe1`	Specify the interface (xe1) to be configured.
`(config-if)#ip address 10.1.1.1/24`	Set the IP address of the interface to 10.1.1.1/24.
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

P

`#configure terminal`	Enter configure mode.
`(config)#interface xe1`	Specify the interface (xe1) to be configured.
`(config-if)#ip address 10.1.1.2/24`	Set the IP address of the interface to 10.1.1.2/24.

(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

Configure IP SLA Configurations on PE 1 router

PE-1

#configure terminal	Enter configure mode.
(config)#ip sla <1-65535>	configure IP SLA with a unique no
(config-ip-sla)# icmp-echo ipv4<destination IP> source-interface <interface name>	configure the icmp-echo using destination Ip Address and source interface name
(config-ip-sla-echo)#threshold <1000-60000>	Configure the threshold value
(config-ip-sla-echo)#timeout <1000-60000>	Configure the Timeout value
(config-ip-sla-echo)#frequency <1-60>	Configure the frequency value
(config-ip-sla-echo)#exit	Exit icmp-echo mode
(config-ip-sla)#exit	Exit from IP SLA mode
(config)#commit	Commit the candidate configuration to the running configuration
(config)#time-range <1-60 characters>	configure a time-range
(config-tr)#start-time 11:22 3 july 2021	configure a start-time
(config-tr)#end-time after 200	Configure end-time
(config-tr)#frequency hourly	configure frequency
(config-tr)#exit	exit from time-range
(config)#ip sla schedule <1-65535> time-range echo_schedule	Schedule a IP SLA measurement
(config)#commit	Commit the candidate configuration to the running configuration

Validation

PE-1

```
#sh running-config ip sla
ip sla 1
 icmp-echo ipv4 10.1.1.2 source-interface xe1
  frequency 6
  threshold 50000
  timeout 55000
ip sla schedule 1 time-range tr1
#sh running-config time-range
!
time-range tr1
 start-time 05:00 21 september 2021
 end-time 06:40 21 september 2021
```

```
#ping 10.1.1.2
Press CTRL+C to exit
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.436 ms
1 packets transmitted, 1 received, 0% packet loss, time 0ms
#sh ip sla summary
IP SLA Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                                   (usec)        Code        Run
-------------------------------------------------------------------
*1      icmp-echo   10.1.1.2         2000         OK        2021 Sep 21 05:01:00

#sh ip sla statistics 1 detail
=======================================
            IP SLA Statistics
=======================================
  IP SLA ID                : 1
  Start Time               : 2021 Sep 21 05:00:00
  Elapsed time(milli sec)  : 25003
  Packets Sent             : 5
  Packets Received         : 5
  Packet Loss(%)           : 0.0000
  Invalid Tests            : 0
  Round Trip Delay(usec)
      Minimum              : 1000
      Maximum              : 1000
      Average              : 800
```

CHAPTER 8 # Traffic Mirroring using ERSPAN

Overview

Encapsulated Remote Switched Port Analyzer (ERSPAN) is a function used for monitoring network traffic. Using ERSPAN, you can mirror traffic from one or more ports or VLANs on a network switch and send the mirrored traffic to a remote monitoring device for analysis.

ERSPAN encapsulates mirrored traffic with Generic Routing Encapsulation (GRE) and, in addition, ERSPAN headers to send over an IP network.

Traffic mirroring protocols such as Switched Port Analyzer (SPAN) and Remote Switched Port Analyzer (RSPAN) in OcNOS allow traffic analysis within the same domain. ERSPAN aims to overcome this limitation by routing the traffic to any destination on the network.

Feature Characteristics

The main characteristics of ERSPAN are as follows:

- Transports mirrored traffic from the source to the destination over Layer 3 IP network.
- Monitors ingress, egress, or both ingress and egress traffic.
- Sends mirrored traffic to remote monitoring device for analysis without being restricted by Layer 2 boundaries.
- Supports filters on ingress traffic providing capability to filter the traffic to be mirrored.
- Supports Type 1 and Type 3 ERSPAN, with Type 1 as the default.

Supported Hardware

- XGS platforms - TD3, TH/TH2

Limitations

The following encapsulations are not supported:

- Type 1 IPv6
- Type 3 IPv4 - egress traffic mirroring
- Type 3 IPv6

Prerequisites

Before configuration, ensure the IP address is available for:

- Destination of the ERSPAN tunnel.
- Origin of the ERSPAN tunnel.

Configuration

The following configuration enables an sender session to send packets to the destination over ERSPAN tunnels.

Topology

The topology shown here consists of Hosts **H1** and **H2**, a Sender node **R1** and a Receiver node **R2.**
The sender node forwards ERSPAN traffic to the receiver node. An ERSPAN tunnel is created between R1 and R2 over interface **ge3**.

R1 collects the traffic received or sent over one or more interfaces (such as **ge1** and/or **ge2)**, mirrors the collected traffic, encapsulates the packets inside ERSPAN and sends them to the IP address on R2.

R2 is configured to receive ERSPAN encapsulated packets.

Figure 8-8: ERSPAN Topology

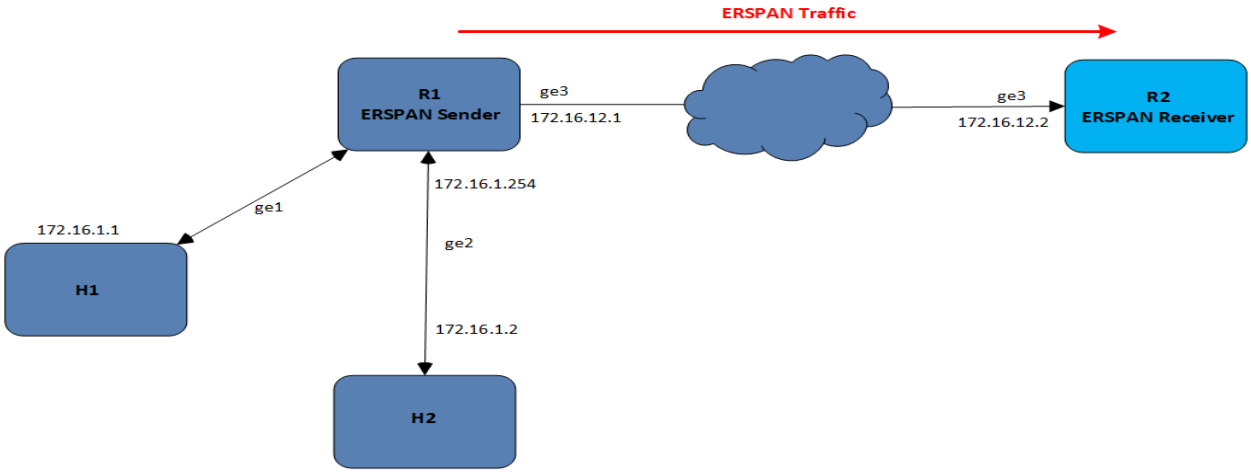

The configuration is done in two stages:

1. Configure ERSPAN destination
2. Configure ERSPAN sender session using the ERSPAN destination

Configure ERSPAN destination

1. Enter configure mode and set a name for the ERSPAN destination.
   ```
   R1(config-router)#monitor destination erspan erspan_dest_1
   R1(config-erspan-dst)#
   ```
2. Configure the destination IPv4/IPv6 where the ERSPAN packets will be forwarded.
   ```
   R1(config-erspan-dst)#dest-ip 172.16.12.2
   ```
3. Set the origin IPv4/IPv6 of the ERSPAN tunnel.
   ```
   R1(config-erspan-dst)#origin-ip 172.16.12.1
   ```
4. The below parameters are optional. If not specified, the default values are used for each parameter.
   - Set the VRF where the ERSPAN tunnel will be created. If not specified, value `default` will be used. `R1(config-erspan-dst)#vrf default`
   - Set the TTL value to be used at the outer IP layer. If not specified, value 255 will be used.

```
R1(config-erspan-dst)#ttl 50
```

- Set the DSCP value to be used at the outer IP layer. If not specified, value 0 will be used.

  ```
  R1(config-erspan-dst)#dscp 211
  ```

- Enable the packet truncation when mirroring to the ERSPAN destination. When this flag is set, the original packet is truncated to 192 bytes and then encapsulated in ERSPAN. By default, truncation is not enabled.

  ```
  R1(config-erspan-dst)#enable-truncate
  ```

  Note: Packet truncation is not supported on TH and TH2 platforms.

- Set the ERSPAN tunnel to Type 1 or Type 3. If not specified, value 1 will be used.

  ```
  R1(config-erspan-dst)#erspan-type 1
  ```

- Set the ERSPAN ID to be used in the ERSPAN session. This is relevant for type 3 only. If not specified, value 0 is used.

  ```
  R1(config-erspan-dst)#erspan-id 100
  ```

- Set a Hardware ID value between 0 to 63. This parameter is relevant for type 3 only. If not specified, value 0 is used.

  ```
  R1(config-erspan-dst)#hardware-id 45
  ```

- Set a Switch ID value between 0 to 511. This parameter is relevant for type 3 only. If not specified, value 0 is used.

  ```
  R1(config-erspan-dst)#switch-id 110
  ```

- Commit the changes.

  ```
  R1(config-erspan-dst)#commit
  ```

Configure ERSPAN sender session

1. Enter configure mode and create a sender session with ID 1. Optionally, you can enter a description for the session (containing a maximum of 32 characters).

   ```
   R1(config)#monitor session 1 type erspan-sender
   R1(config-monitor)#description R1 ERSPAN sender
   ```

2. Configure the ERSPAN destination for the session using the name of the destination that has been created previously.

   ```
   R1(config-monitor)#destination erspan erspan_dest_1
   ```

3. Optionally, add sources such as `source VLAN` and/or `source interface` to the sessions. For example, the command `source interface` configures the monitored source interface and the direction of the traffic to be monitored. If not specified, both ingress and egress traffic are monitored.

   ```
   R1(config-monitor)#source interface ce51 rx
   ```

4. Enable the configured session on the interface.

   ```
   no shut
   ```

ERSPAN Snippet Configuration

To verify the configuration and view the overall commands, use the `show running-config monitor` command.

```
R1#show running-config monitor
monitor destination erspan erspan_dest_1
  dest-ip 23.1.1.2
  vrf default
  origin-ip 69.69.69.69
  ttl 211
  dscp 50
  enable-truncate
  erspan-type 1
```

```
!
monitor session 1 type erspan-sender
 description R1 ERSPAN sender
 source interface ce51 rx
 destination erspan erspan_dest_1
 no shut
```

Validation

To verify the ERSPAN configuration, check the output of the `show monitor session 1` command.

```
#show monitor session 1
session 1
---------------
description : R1 ERSPAN sender
type : ERSPAN Sender
state : up
source intf :
tx :
rx : ge1
both :
source VLANs :
rx :
destination ERSPAN: erspan_dest_1
ERSPAN Type : 1
Dest IP addr : 172.16.12.2
Origin IP addr: 172.16.12.1
Dest VRF : default
ERSPAN ID : 0
DSCP : 50
TTL : 211
pkt truncate : Enabled
NextHop addr : 172.16.12.2
NextHop intf : ge3
filter count :
Legend: f = forwarding enabled, l = learning enabled
Sender#
```

CLI Commands

The ERSPAN feature introduces the following configuration commands.

destination ERSPAN

Use this command to configure the ERSPAN destination for an ERSPAN sender session. The destination must be already created using the command `monitor destination erspan`.

Use `no` form of this command to remove the ERSPAN destination from the session.

Command Syntax

```
destination erspan NAME
```

```
no destination erspan
```

Parameters

NAME	ERSPAN destination name mentioned in the command `monitor destination erspan`

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to configure the ERSPAN destination.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#dest-ip 172.16.12.2
(config-erspan-dst)#origin-ip 172.16.12.1
(config-erspan-dst)#exit
(config)#monitor session 1 type erspan-sender
(config-monitor)#destination erspan erspan_dest_1
(config-monitor)#no destination erspan
```

ERSPAN origin ip

Use this command to set the origin IPv4/IPv6 of the ERSPAN tunnel.

Use `no` form of this command to unset the origin IPv4/IPv6 of the ERSPAN tunnel.

Command Syntax

```
origin-ip A.B.C.D/X:X::X:X
no origin-ip
```

Parameters

A.B.C.D/X:X::X:X	Origin IPv4/IPv6 address of the ERSPAN tunnel

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the origin IPv4/IPv6 of the ERSPAN tunnel.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#origin-ip 172.16.12.1
(config-erspan-dst)#commit

(config-erspan-dst)#no origin-ip
(config-erspan-dst)#commit
```

ERSPAN destination ip

Use this command to set the destination IPv4/IPv6 of the ERSPAN tunnel.

Use `no` form of this command to unset the destination IPv4/IPv6 of the ERSPAN tunnel.

Command Syntax

```
dest-ip A.B.C.D/X:X::X:X
no dest-ip
```

Parameters

`A.B.C.D/X:X::X:X`	Destination IPv4/IPv6 address of the ERSPAN tunnel

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the destination IPv4/IPv6 of the ERSPAN tunnel.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#dest-ip 172.16.12.1
(config-erspan-dst)#commit

(config-erspan-dst)#no dest-ip
(config-erspan-dst)#commit
```

ERSPAN vrf

Use this command to set the VRF where the ERSPAN tunnel will be

created. Use `no` form of this command to reset the VRF to default.

Command Syntax

```
vrf VRF_NAME
no vrf
```

Parameters

VRF_NAME	VRF name where the ERSPAN tunnel will be created

Default

Default

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the VRF where the ERSPAN tunnel will be created.

```
((config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#vrf custom_vrf_1
(config-erspan-dst)#commit

(config-erspan-dst)#no vrf
(config-erspan-dst)#commit
```

ERSPAN ip ttl

Use this command to set the Time To Live (TTL) value to use at the outer IP layer. This is an optional parameter that uses TTL value 255, if not specified.

Use `no` form of this command to reset the TTL value to 255.

Command Syntax

```
ttl <1-255>
```

```
no ttl
```

Parameters

<1-255>	TTL value to be used

Default

Value 255

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the TTL value to use at the outer IP layer.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#ttl 25
(config-erspan-dst)#commit

(config-erspan-dst)#no ttl
(config-erspan-dst)#commit
```

ERSPAN ip dscp

Use this command to set the Differentiated Services Code Point (DSCP) value to use at the outer IP layer. This is an optional parameter that uses DSCP value 0, if not specified.

Use `no` form of this command to reset the DSCP value to 0.

Command Syntax

```
dscp <0-63>
no dscp
```

Parameters

`<0-63>`	DSCP value to be used

Default

Value 0

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the DSCP value to use at the outer IP layer.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#dscp 42
(config-erspan-dst)#commit

(config-erspan-dst)#no dscp
(config-erspan-dst)#commit
```

ERSPAN enable truncate

Use this command to enable packet truncation when mirroring to the ERSPAN destination. When this flag is set, the original packet is truncated to 192 bytes and then encapsulated in ERSPAN.

Use `no` form of this command to disable packet truncate.

Command Syntax

```
enable-truncate
no enable-truncate
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to enable the packet truncation.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#enable-truncate
(config-erspan-dst)#commit

(config-erspan-dst)#no enable-truncate
(config-erspan-dst)#commit
```

ERSPAN type

Use this command to set the ERSPAN tunnel to Type 1 or Type 3. Note that ERSPAN Type 2 is not supported on XGS TD3 and TH/TH2 boards.

Use `no` form of this command to reset the ERSPAN type to the default value.

Command Syntax

```
erspan-type (1|3)
no erspan-type
```

Parameters

`1`	Use ERSPAN Type 1
`3`	Use ERSPAN Type 3

Default

Type 1

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the ERSPAN tunnel.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#erspan-type 3
(config-erspan-dst)#commit

(config-erspan-dst)#no erspan-type
(config-erspan-dst)#commit
```

ERSPAN id

Use this command to set the ERSPAN ID to be used in the ERSPAN session. This is only relevant for ERSPAN Type 3. This is an optional parameter and the ERSPAN ID 0 is used, if not specified.

Use `no` form of this command to reset the value to 0.

Command Syntax

```
erspan-id (1-1023)
no erspan-id
```

Parameters

`<1-1023>`	ERSPAN ID to be used

Default

Value 0

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the ERSPAN ID.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#erspan-id 33
(config-erspan-dst)#commit

(config-erspan-dst)#no erspan-id
(config-erspan-dst)#commit
```

ERSPAN hardware id

Use this command to set the Hardware ID to be used. This is only relevant for ERSPAN Type 3.

Use `no` form of this command to reset the value to 0.

Command Syntax

```
hardware-id (0-63)
no hardware-id
```

Parameters

<1-63>	Hardware ID to be used

Default

Value 0

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the Hardware ID.

```
(config)#monitor destination erspan erspan_dest_1
(config-erspan-dst)#hardware-id 12
(config-erspan-dst)#commit

(config-erspan-dst)#no hardware-id
(config-erspan-dst)#commit
```

ERSPAN switch id

Use this command to set value for the Switch ID to be used. This is only relevant for ERSPAN Type 3.

Use `no` form of this command to reset the value to 0.

Command Syntax

```
switch-id (0-1023)
no switch-id
```

Parameters

<1-1023>	Switch ID to be used

Default

Value 0

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

The following sequence of commands is used to set the Switch ID.

```
(config)#monitor destination erspan erspan_dest_1
```

```
(config-erspan-dst)#switch-id 112
(config-erspan-dst)#commit

(config-erspan-dst)#no switch-id
(config-erspan-dst)#commit
```

The below commands have been revised for ERSPAN. For more details, refer to the *Traffic Mirroring Commands* chapter.

- `monitor session`
- `show monitor session` - Example section

Glossary

The following table provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	**Description**
Switched Port Analyzer (SPAN)	A protocol that monitors the traffic on source port and sends a copy of the traffic to a destination port.
Remote Switched Port Analyzer (RSPAN)	A protocol that monitors the traffic distributed over multiple switches from the source ports.
Time to Live (TTL)	A limit on how long a piece of information can exist before it should be discarded.
Differentiated Services Code Point (DSCP)	A six-bit field in an IP header that enables allocation of resources on a per-packet basis.
Virtual Routing and Forwarding (VRF)	A technology that allows multiple data structures to co-exist within the same router at the same time.

Fundamental Layer 3 Command Reference

© 2025 IP Infusion Inc. Proprietary

CHAPTER 1 Fundamental Layer 3 Commands

This chapter describes the fundamental Layer 3 commands:

- automatic-router-id-selection enable
- clear ip route kernel
- clear ip route
- clear ip route vrf NAME
- clear router-id
- debug rib
- description
- fib retain
- ip route
- ip urpf enable
- ip verify unicast source reachable-via
- ip vrf
- ipv6 route
- maximum-paths
- max-static-routes
- router-id
- show debugging rib
- show ip rpf
- show ipv6 rpf
- show router-id
- show running-config router
- show running-config router-id
- show running-config urpf
- show running-config vrf
- snmp restart rib

© 2025 IP Infusion Inc. Proprietary

automatic-router-id-selection enable

Use this command to assure that OcNOS selects the loopback IP address as the router-id each time the device is rebooted.

Use the `no` form of this command to remove this constraint.

Command Syntax

```
automatic-router-id-selection enable
```

```
no automatic-router-id-selection enable
```

Parameters

None

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#automatic-router-id-selection enable
(config)#
```

clear ip route kernel

Use this command to clear stale IPv4 routes from the RIB (Routing Information Base) and FIB (Forwarding Information Base).

Command Syntax

```
clear ip route kernel
clear ip kernel route
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip route kernel
```

clear ip route

Use this command to clear an all IPv4 routes or any specific prefix routes.

Command Syntax

```
clear ip route (*|A.B.C.D/M)
```

Parameters

`*`	Clears all routes
`A.B.C.D/M`	Prefix to be cleared

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip route *
#clear ip route 1.1.1.0/24
```

clear ip route vrf NAME

Use this command to clear all IPv4 VRF route or any specific prefix VRF route of any particular VRF name.

Command Syntax

```
clear ip route vrf NAME (*|A.B.C.D/M)
```

Parameters

NAME	VPN Routing or Forwarding instance name
*	Clears all routes
A.B.C.D/M	Prefix to be cleared

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip route vrf myVRF *
```

clear router-id

Use this command to clear the current Router-Id and trigger the Router-Id calculation again. The new Router-id is sent to all protocol modules.

- To clear only the router-id for the default VRF, enter `clear router-id`.
- To clear the router-id for a particular VRF, enter `clear router-id vrf VRFNAME`.
- To clear the router-id for all VRFs, enter `clear router-id vrf`.

Command Syntax

```
clear router-id (vrf (VRFNAME|))
```

Parameters

`VRFNAME`	VPN routing/forwarding instance name.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear router-id
#clear router-id vrf
```

debug rib

Use this command to debug the `ribd` process.

Use the `no` form of this command to stop debugging.

Command Syntax

```
debug rib (all|)
debug rib events
debug rib packet (recv|send|) (detail|)
debug rib nsm
debug rib bfd
debug rib kernel
debug rib monitor
debug ip routing (add-route|delete-route|mod-route|
no debug all
no debug rib (all|)
no debug all rib
no debug rib events
no debug rib packet (recv|send|) (detail|)
no debug rib nsm
no debug rib bfd
no debug rib kernel
no debug rib monitor
no debug ip routing (add-route|delete-route|mod-route|)
```

Parameters

`all`	All debugging functions
`events`	Events
`packet`	Packet events
`recv`	Received packets
`send`	Sent packets
`detail`	Detailed information
`nsm`	NSM events
`kernel`	RIB kernel
`monitor`	Enable Monitor route netlink
`bfd`	BFD (Bidirectional Forwarding Detection) events
`ip routing`	IPv4 routing events
`add-route`	Add route events

`delete-route`
: Delete route events

`mod-route`
: Modify route events

Disabled

By default, debug command is disabled.

Command Mode

Privileged Exec mod

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#debug rib all
```

description

Use this command to assign a description to a Virtual Router instance.

Use the `no` parameter to remove a description.

Command Syntax

```
description LINE
no description
```

Parameters

`LINE`	Virtual Router description maximum 80 characters

Disabled

By default, description command is disabled

Command Mode

VR mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#virtual-router VR1
(config-vr)#description VR1 has been created for CLI testing
(config-vr)#exit

(config)#virtual-router VR1
(config-vr)#no description
(config-vr)#exit
```

fib retain

Use this command to set the retention time for stale routes in the Forwarding Information Base (FIB) when `ribd` restarts. The `ribd` process reads the FIB and treats previously self-installed routes as stale.

You can display stale routes by running the `show ip route database` command. All routes preceded by the symbol `p` are stale routes. When protocol modules restart, `ribd` overrides these stale routes with routes updated by the protocol modules.

Table 8 show the behavior of routes when `ribd` stops.

Table 8: FIB retention

Command	Behavior
`fib retain`	Does not clear routes from the FIB and retains stale routes for 60 seconds when restarted.
`fib retain forever`	Does not clear routes and retains stale routes forever.
`fib retain time <1-65535>`	Does not clear routes and retains stale routes for the specified seconds.
`no fib retain` (default)	Cleans up routes in the FIB, but retains stale routes for 60 seconds when restarted.

You can remove stale routes at any time with the clear ip route kernel command.

Use the `no` form of this command to revert to default; that is, do not retain routes in the FIB when `ribd` stops.

Command Syntax

```
fib retain (forever|time <1-65535>|)
no fib retain (forever|time <1-65535>|)
```

Parameters

`forever`	Retain FIB forever
`time`	Retain FIB for a time after `ribd` restarts
`<1-65535>`	Retention time in seconds; if you omit this value, the default is 60 seconds

Default

Routes are cleared from the FIB when `ribd` stops. However, when `ribd` restarts, stale routes are retained for 60 seconds.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#fib retain time 180
```

ip route

Use this command to create an IPv4 static route.

Use the `no` form of this command to delete a static route.

Note: Equal-Cost Multi-Path (ECMP) static routes with one next-hop residing in a local Virtual Routing and Forwarding (VRF) instance while the other next-hop is in an inter-VRF are not supported. Additionally, ECMP static routes within an Inter-VRF (IVRF) are not supported.

Command Syntax

```
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME (<1-255>|) (tag <0-
  4294967295>|) (track <1-500>|) (description LINE|)
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D (<1-255>|) (tag <0-
  4294967295>|) (track <1-500>|) (recursive|) (description LINE|)
ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME (<1-255>|) (tag <0-
  4294967295>|) (track <1-500>|) (recursive|) (description LINE|)
ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME global (track <1-
  500>|) (recursive|)
ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME global (track <1-500>|)
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME {<1-255>|tag <0-
  4294967295>|track <1-500>|description LINE}
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D {<1-255>|tag <0-
  4294967295>|track <1-500>|recursive|description LINE}
no ip route (vrf NAME|) (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME {<1-255>|tag <0-
  4294967295>|track <1-500>|recursive|description LINE}
no ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME global
no ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) A.B.C.D IFNAME global {track <1-
  500>|recursive}
no ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME global
no ip route vrf NAME (A.B.C.D/M|A.B.C.D A.B.C.D) IFNAME global (track <1-500>|)
```

Parameters

`vrf NAME`	Specifies the user-defined VRF instance name.
`A.B.C.D/M`	Denotes the subnet by specifying the IP address destination prefix along with a mask length.
`A.B.C.D A.B.C.D`	

	Specifies the subnet with an IP destination address and its corresponding mask.
`A.B.C.D`	Specifies the gateway's next-hop IPv4 address.
`IFNAME`	Specifies the gateway's next-hop interface name.
`global`	Global table lookup (to support inter-VRF static route leaking).
`<1-255>`	Specifies the administrative distance value.
`track <1-500>`	Associates a tracking ID with an IPv4 static route for the tracked object.
`description LINE`	Adds a description to the static route, with a maximum character limit of 80.
`tag <0-4294967295>`	Assigns a tag value to the route, which can be used as a “match” value to control route redistribution through route maps.
`recursive`	Enables recursive lookup behavior for the next-hop in each static route. By default, it is disabled.

Default

None

Command Mode

Configure mode

Applicability

The command was introduced before OcNOS version 1.3 and was updated in OcNOS version 1.3.4. Introduced new parameters `track <1-500>` in OcNOS version 5.1 and `recursive` in OcNOS version 6.6.0.

Examples

```
#configure terminal
(config)#ip route 192.168.3.0 255.255.255.0 2.2.2.2 128
(config)#ip route 1.1.1.0/24 eth0 32
(config)#ip route vrf new 1.1.1.1/1 1.1.1.1 eth1 description new tag 1
(config)# ip route 40.1.1.0/24 eth1 track 10
(config)#commit

(config)#no ip route 40.1.1.0/24 eth1 track 10
(config)#commit
```

This example creates VRF static routes with the nexthops belonging to the default VRF. The nexthop gateway address can be the IFNAME network address or any other IP address reachable via IFNAME.

```
#show ip route vrf
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     10.12.19.0/24 is directly connected, eth0, 4d23h06m
C     11.1.1.0/24 is directly connected, eth4, 4d23h01m
```

```
C     13.13.13.0/31 is directly connected, eth3, 4d23h06m
C     50.5.5.0/24 is directly connected, eth3, 4d23h06m
C     100.100.100.0/24 is directly connected, lo, 4d23h06m
O     101.1.1.1/32 [110/11] via 11.1.1.2, eth4, 19:20:39
C     127.0.0.0/8 is directly connected, lo, 4d23h06m
Gateway of last resort is not set

(config)#ip route vrf vrf1 201.201.201.201/32 11.1.1.11 eth4 global
(config)#ip route vrf vrf1 202.202.202.202/32 101.1.1.1 eth4 global

#show ip route vrf vrf1 static
IP Route Table for VRF "vrf1"
S    v201.201.201.201/32 [1/0] via 11.1.1.11, eth4, 00:00:44
S    v202.202.202.202/32 [1/0] via 101.1.1.1, eth4 (recursive via 11.1.1.2), 00:00:17
Gateway of last resort is not set

#configure terminal
(config)#ip route vrf vrf1 20.2.2.0/24 xe1 global
(config)#commit

(config)#no ip route vrf vrf1 20.2.2.0/24 xe1 global track 10
(config)#commit
```

ip urpf enable

Use this command to enable uRPF mode on the system.

Use the `no` form of the command to disable uRPF mode on this system.

Note: The configuration is applied only after a reboot.

Command Syntax

```
ip urpf enable
no ip urpf enable
```

Parameter

None

Default

By default, uRPF mode on the system is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip urpf enable
% System Reboot is required for new URPF configurations to take effect

(config)#no ip urpf enable
% System Reboot is required for new URPF configurations to take effect
```

ip verify unicast source reachable-via

Use this command to Configure Unicast Reverse Path Forwarding with source-specific reachability on the interface level.

Use the no form of the command to disable Unicast Reverse Path Forwarding with source-specific reachability on the interface level.

Command Syntax

```
ip verify unicast source reachable-via (any (allow-default|) | rx )
no ip verify unicast source reachable-via
```

Parameters

`any`	Source is reachable via any interface
`rx`	Source is reachable via an interface on which packet was received
`allow-default`	Loose Default Route Unicast Reverse Path Forwarding

Default

N/A

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS Version 6.0.0.

Examples

```
#configure terminal
(config)#interface xe1
(config-if)#ip verify unicast source reachable-via any
(config-if)#ip verify unicast source reachable-via any allow-default
(config-if)#ip verify unicast source reachable-via rx
(config-if)#no ip verify unicast source reachable-via
```

ip vrf

This command creates a user-defined VRF (Virtual Routing and Forwarding) RIB (Routing Information Base), assigns a VRF identifier, and switches to VRF mode.

Use the `no` parameter with command to remove a VRF RIB.

Note: On configuration of user-defined VRF with `ip vrf WORD` CLI, `lo.WORD` interface configuration will not be displayed in `show running-config` even though `lo.WORD` interface, default IPv4 address and default IPv6 address are present in system. To display the `interface lo.WORD`, `ip address 127.0.0.1/8` and `ipv6 address ::1/128` configurations in show running configuration user needs to explicitly configure the same from CLI. This is due to a restriction with scaled VRF configurations in a single commit.

Command Syntax

```
ip vrf WORD
no ip vrf WORD
```

Parameter

`WORD`	VRF identifier

Default

By default, no user-defined VRFs exist, only the default VRF.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#ip vrf myVRF
(config-vrf)#
```

ipv6 route

Use this command to create an IPv6 static route.

Use the no form of this command to delete a static route

Note: Using an interface name as the gateway next-hop for an IPv6 static route is not supported.

Command Syntax

```
ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME

ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X

ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME

ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME (<1-255>|) (track <1-
  500>|) (description LINE|)

ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X (<1-255>|) (track
  <1-500>|) (recursive|) (description LINE|)

ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME (<1-255>|)
  (track <1-500>|) (recursive|) (description LINE|)

ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME global (track
  <1-500>|) (recursive|)

ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME global (track <1-500>|)

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME {<1-255>|track <1-
  500>|description LINE}

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X {<1-255>|track
  <1-500>|recursive|description LINE}

no ipv6 route (vrf NAME|) (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME {<1-
  255>|track <1-500>|recursive|description LINE}

no ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME global

no ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) X:X::X:X IFNAME global {track
  <1-500>|recursive}

no ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME global

no ipv6 route vrf NAME (X:X::X:X/M|X:X::X:X X:X::X:X) IFNAME global (track <1-
  500>|)
```

Parameters

`vrf NAME`	Specifies the user-defined VRF instance name.
`X:X::X:X/M`	Denotes the subnet by specifying the IPv6 address destination prefix along with a mask length.
`X:X::X:X X:X::X:X`	Specifies the subnet with an IPv6 destination address and its corresponding mask.
`X:X::X:X`	Specifies the gateway's next-hop IPv6 address.

IFNAME	Specifies the gateway's next-hop interface name.
global	Global table lookup (to support inter-VRF static route leaking).
<1-255>	Specifies the administrative distance value.
track <1-500>	Associates a tracking ID with an IPv6 static route for the tracked object.
description LINE	Adds a description to the static route, with a maximum character limit of 80.
recursive	Enables recursive lookup behavior for the next-hop in each static route. By default, it is disabled.

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 1.3.

Example

```
#configure terminal
(config)#ipv6 route 1234::/64 1000::1 128
(config)#ipv6 route 1235::/64 eth1 32
(config)#ipv6 route vrf vrf3 1236::/64 eth3
(config)#end

#show ipv6 route vrf all
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
IA - OSPF inter area, E1 - OSPF external type 1,
E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime
IP Route Table for VRF "default"
C ::1/128 via ::, lo, 23:35:55
C 2::2/128 via ::, lo, 23:35:55
C 1000::/64 via ::, eth1, 00:07:45
S 1234::/64 [128/0] via 1000::1, eth1, 00:07:04
S 1235::/64 [32/0] via ::, eth1, 00:06:38
C 6000::/64 via ::, eth1, 23:35:55
C 7000::/64 via ::, eth7, 23:35:55
C fe80::/64 via ::, eth7, 23:35:55
IP Route Table for VRF "management"
IP Route Table for VRF "vrf3"
S 1236::/64 [1/0] via ::, eth3, 00:00:13
C 3000::/64 via ::, eth3, 00:07:55
C fe80::/64 via ::, eth3, 00:08:05

#show ipv6 route vrf all static
IP Route Table for VRF "default"
```

```
S 1234::/64 [128/0] via 1000::1, eth1, 00:15:28
S 1235::/64 [32/0] via ::, eth1, 00:15:02
IP Route Table for VRF "management"
IP Route Table for VRF "vrf3"
S 1236::/64 [1/0] via ::, eth3, 00:08:37
```

maximum-paths

Use this command to set the maximum number of paths to install in the FIB (Forwarding Information Base) for the ECMP (Equal-Cost MultiPath) feature.

Use the `no` parameter with this command to revert to default.

Note: If you change the number of paths from the default (8), you must save the running configuration and perform a reboot.

Command Syntax

```
maximum-paths <1-64>
no maximum-paths
```

Parameter

`<1-64>`	Maximum number of paths to install in the FIB

Default

By default, the maximum number of paths is 8.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#maximum-paths 5
```

max-static-routes

Use this command to set the maximum number of static routes.

Use the `no` parameter to disable this command.

Command Syntax

```
max-static-routes <1-4294967294>
no max-static-routes
```

Parameters

`<1-4294967294>` Maximum number of static routes

Default

By default, max static routes value is 4294967294

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#max-static-routes 123

(config)#no max-static-routes
```

router-id

Use this command to add a router identifier for this system.

Use the `no` form of this command to disable this function.

Command Syntax

```
router-id A.B.C.D
no router-id (A.B.C.D)
```

Parameters

`A.B.C.D`	Router identifier in IP address format for this system.

Default

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router-id 123.12.3.123
(config)#
```

show debugging rib

Use this command to display debug settings.

Command Syntax

```
show debugging rib
```

Parameters

None

Command Mode

Privileged Exec Mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show debugging rib
```

show ip rpf

Use this command to display reverse path forwarding (RPF) information for the specified source address.

Command Syntax

```
show ip rpf A.B.C.D
show ip rpf (vrf NAME|) A.B.C.D
```

Parameters

A.B.C.D	IP address of multicast source.
NAME	Virtual Routing and Forwarding name.

Default

None

Command Mode

Exec and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip rpf 10.10.10.50

RPF information for 10.10.10.50
  RPF interface: eth0
  RPF neighbor: 10.1.2.1
  RPF route: 0.0.0.0/0
  RPF type: unicast (kernel)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
  Distance: 0
  Metric: 0
#
```

Table 9 explains the output fields.

Table 9: show ip rpf output fields

Field	Description
RPF Interface	Name of the RPF interface.
RPF neighbor	Upstream RPF neighbor.
RPF route	Route table in which the logical interface address is located.
RPF type	Different type of RPF like multicast, unicast, MBGP, DVMRP, or static mroutes.

Table 9: show ip rpf output fields

Field	Description
RPF recursion count	Number of times that the router lookups its routing table more than once to find out the immediate next-hop and exiting interface.
Distance	IP address of the remote side of the connection. Doing distance-preferred lookups across tables.
Metric	Metrics are informational units that can be measured and compared.

show ipv6 rpf

Use this command to display reverse path forwarding (RPF) information for the specified source address.

Command Syntax

```
show ipv6 rpf X:X::X:X
show ipv6 rpf (vrf NAME|) X:X::X:X
```

Parameters

X:X::X:X	IP address of multicast source.
NAME	Virtual Routing and Forwarding name.

Default

None

Command Mode

Exec and privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ipv6 rpf 10:10::10:50

RPF information for 10.10.10.50
  RPF interface: eth0
  RPF neighbor: 10.1.2.1
  RPF route: 0.0.0.0/0
  RPF type: unicast (kernel)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
  Distance: 0
  Metric: 0
#
```

Table 10 explains the output fields.

Table 10: show ipv6 rpf output fields

Field	Description
RPF Interface	Name of the RPF interface.
RPF neighbor	Upstream RPF neighbor.
RPF route	Route table in which the logical interface address is located.
RPF type	Different type of RPF like multicast, unicast, MBGP, DVMRP, or static mroutes.

Table 10: show ipv6 rpf output fields

Field	Description
RPF recursion count	Number of times that the router lookups its routing table more than once to find out the immediate next-hop and exiting interface.
Distance	IPv6 address of the remote side of the connection. Doing distance-preferred lookups across tables.
Metric	Metrics are informational units that can be measured and compared.

show router-id

Use this command to display the Router ID of the current system.

Command Syntax

```
show router-id
```

Parameters

None

Default

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show router-id
Router ID: 10.55.0.2 (automatic)
```

show running-config router

Use this command to display the running system router configuration.

Command Syntax

```
show running-config router bgp
show running-config router isis
show running-config router ospf
show running-config router rip
show running-config router vrrp
```

Parameters

`bgp`	Display Border Gateway Protocol (BGP) information.
`isis`	Display Intermediate System to Intermediate System (IS-IS) information.
`ospf`	Display Open Shortest Path First (OSPF) information.
`rip`	Display Routing Information Protocol (RIP) information.
`vrrp`	Display Virtual Router Redundancy Protocol (VRRP) information.

Default

None

Command Mode

Privileged exec mode, configure mode, router-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#show running-config router vrrp
!
router-id 3.3.3.3
!
```

show running-config router-id

Use this command to show the running system global router ID configuration.

Command Syntax

```
show running-config router-id
```

Parameters

None

Default

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#show running-config router-id
!
router-id 3.3.3.3
!
```

show running-config urpf

Use this command to check uRPF status for this system.

Command Syntax

```
show running-config urpf
```

Parameters

None

Default

None

Command Mode

Exec mode and Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show running-config urpf
ip urpf enable

(config)#show running-config urpf
ip urpf enable
```

show running-config vrf

Use this command to show the running system status and configuration details for a specified VRF instance name.

Command Syntax

```
show running-config vrf WORD
```

Parameters

WORD	Virtual Routing and Forwarding name

Command Mode

Privileged exec mode, configure mode, router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#show running-config vrf xyz
!
ip vrf xyz
 description vrf
 router-id 11.11.11.11
```

snmp restart rib

Use this command to restart SNMP in Routing Information Base (RIB)

Command Syntax

```
snmp restart rib
```

Parameters

None

Default

By default, snmp restart command is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#config terminal
(config)#snmp restart rib
```

CHAPTER 2 Route-Map Commands

This chapter describes route-map commands.

- continue
- match as-path
- match community
- match extcommunity
- match interface
- match ip address
- match ip address prefix-list
- match ip next-hop
- match ip next-hop prefix-list
- match ip peer
- match ipv6 address
- match ipv6 address prefix-list
- match ipv6 next-hop
- match ipv6 next-hop prefix-list
- match ipv6 peer
- match metric
- match origin
- match route-type
- match tag
- route-map
- set aggregator
- set as-path
- set atomic-aggregate
- set comm-list
- set community
- set dampening
- set extcommunity
- set interface null0
- set ip next-hop
- set ipv6 next-hop
- set level
- set local-preference
- set metric
- set metric-type
- set origin

© 2025 IP Infusion Inc. Proprietary

- set originator-id
- set tag
- set weight
- show route-map
- show running-config route-map

continue

The continue clause provides the capability to execute additional entries in a route map after an entry is executed with a successful match and set clauses (i.e), the continue command allows multiple entries to be evaluated within a single route-map. Continue commands can be assigned optional sequence numbers that indicate the order in which clauses are to be evaluated.

Use the no form of the command (specifying a particular sequence number if desired), to remove individual or all continue clauses from a route-map.

Continue clause with match Clauses

When a match clause exists in a route map with continue clause, then the continue clause is executed only when a successful match occurs. If a match clause does not exist in the route-map and if a continue clause does exist, the continue clause will be evaluated and then go to the specified route-map entry. When a successful match occurs and we have a continue clause, the route-map executes the set clauses and then goes to the specified route-map entry. If continue clause does not exist in the next route map, then the route-map will behave normally. If a continue clause exists in the next route-map but a match is not successful, the route-map will not continue and will *fall through* to the next sequence number if one exists.

Continue clause with set Clauses

Set clauses are executed after the route-map evaluation is done. The set clauses are evaluated and executed in the order in which they were configured. Set clauses are only executed after a successful match occurs. The continue statement proceeds to the specified route-map entry only after configured set actions are performed. If a set action is configured in the first route-map and then the same set action occurs again, but with a different value in a subsequent route-map entry, then the last set action will override the previous set actions, which were configured with the same set command.

Command Syntax

```
continue <2-65535>|)
no continue <2-65535>|)
```

Parameter

<2-65535>	Continue sequence number.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#route-map Rmap1
(config-route-map)#continue ?
  <2-65535>  Route-map entry sequence number
  <cr>
(config-route-map)#continue 10
```

```
(config-route-map)#continue 30
```

match as-path

Use this command to match an autonomous system path access list. This command specifies the autonomous system path to be matched. If there is a match for the specified AS path, and `permit` is specified, the route is redistributed or controlled, as specified by the set action. If the match criteria are met, and `deny` is specified, the route is not redistributed or controlled. If the match criteria are `not` met then the route is neither accepted nor forwarded, irrespective of `permit` or `deny` specifications.

The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes, depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid only for BGP.

Use the `no` parameter with this command to remove a path list entry.

Command Syntax

```
match as-path WORD
no match as-path (WORD|)
```

Parameter

`WORD`	Autonomous system path access list name.

Default

Enabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map myroute deny 34
(config-route-map)#match as-path myaccesslist
```

match community

Use this command to specify the community to be matched.

Communities are used to group and filter routes. They are designed to provide the ability to apply policies to large numbers of routes by using match and set commands. Community lists are used to identify and filter routes by their common attributes. This command allows the matching based on community lists.

The values set by the `match community` command overrides the global values. The route that does not match at least one match clause is ignored.

Note: This command is valid only for BGP.

Use the `no` parameter with this command to remove the community list entry.

Command Syntax

```
match community (<1-99>|<100-199>|WORD) (exact-match|)
no match community (<1-99>|<100-199>|WORD|) (exact-match|)
```

Parameters

`<1-99>`	Community-list number (standard).
`<100-199>`	Community-list number (expanded).
`WORD`	Community-list name.
`exact-match`	Do exact matching of communities.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#match community mylist
```

match extcommunity

Use this command to match BGP external community list

Communities are used to group and filter routes. They are designed to provide the ability to apply policies to large numbers of routes by using match and set commands. Community lists are used to identify and filter routes by their common attributes. This command allows the matching based on community lists.

The values set by this command overrides the global values. The route that does not match at least one match clause is ignored.

Note: This command is valid only for BGP.

Use the `no` parameter with this command to remove the community list entry.

Command Syntax

```
match extcommunity (<1-99>|<100-199>|WORD) (exact-match|)
no match extcommunity (<1-99>|<100-199>|WORD|) (exact-match|)
```

Parameters

`<1-99>`	Community-list number (standard).
`<100-199>`	Community-list number (expanded).
`WORD`	Name of the community-list.
`exact-match`	Do exact matching of communities.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#match extcommunity mylist
```

match interface

Use this command to define the interface match criterion. This command specifies the next-hop interface name of a route to be matched.

Use the `no` parameter with this command to remove the specified match criterion.

Command Syntax

```
match interface IFNAME
no match interface (IFNAME|)
```

Parameter

`IFNAME`	Interface name.

Default

By default, match interface is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map mymap1 permit 10
(config-route-map)#match interface eth0
```

match ip address

Use this command to specify the match address of route. If there is a match for the specified IP address, and `permit` is specified, the route is redistributed or controlled, as specified by the set action. If the match criteria are met, and `deny` is specified then the route is not redistributed or controlled. If the match criteria are not met, the route is neither accepted nor forwarded, irrespective of `permit` or `deny` specifications.

The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes, depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Use the `no` parameter with this command to remove the `match ip address` entry.

Command Syntax

```
match ip address (<1-199>|<1300-2699>|WORD)
no match ip address (<1-199>|<1300-2699>|WORD|)
```

Parameters

`WORD`	IP access-list name.
`<1-199>`	IP access-list number (standard range).
`<1300-2699>`	IP access-list number (expanded range).

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#match ip address List1
```

match ip address prefix-list

Use this command to match entries of a prefix-list. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Use the `no` parameter with this command too disable this function

Command Syntax

```
match ip address prefix-list WORD
no match ip address prefix-list (WORD|)
```

Parameter

`WORD`	IP prefix list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match ip address prefix-list mylist
```

match ip next-hop

Use this command to specify a next-hop address to be matched in a route-map. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
match ip next-hop (<1-199>|<1300-2699>|WORD)
no match ip next-hop (<1-199>|<1300-2699>|WORD|)
```

Parameters

`WORD`	Specify the IP access-list name.
`<1-199>`	Specify the IP access-list number (standard range).
`<1300-2699>`	Specify the IP access-list number (expanded range).

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match ip next-hop mylist
```

match ip next-hop prefix-list

Use this command to specify the next-hop IP address match criterion using the prefix-list. This command matches the next-hop IP address of a route.

Use the `no` parameter with this command to remove the specified match criterion.

Command Syntax

```
match ip next-hop prefix-list WORD
no match ip next-hop prefix-list (WORD|)
```

Parameter

`WORD`	Prefix-list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#match ip next-hop prefix-list list1
```

match ip peer

Use this command to specify the match peer IPv4 address of a route.

Use the `no` parameter with this command to remove the specified match criterion.

Command Syntax

```
match ip peer (<1-199>|<1300-2699>|WORD)
no match ip peer (<1-199>|<1300-2699>|WORD|)
```

Parameter

`WORD`	IP access-list name.
`<1-199>`	IP access-list number (standard range).
`<1300-2699>`	IP access-list number (expanded range).

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#match ip peer 123

(config-route-map)#no match ip peer 123
```

match ipv6 address

Use this command to specify the match address of route. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Use the `no` parameter with this command to remove the `match ip address` entry.

Note: This command is valid for BGP, OSPFv3, and RIPng only.

Command Syntax

```
match ipv6 address WORD
no match ipv6 address (WORD|)
```

Parameter

`WORD`	IPv6 access list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myRM deny 1
(config-route-map)#match ipv6 address myRM
```

match ipv6 address prefix-list

Use this command to match entries of a prefix-list. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes, depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid for BGP, OSPFv3, and RIPng only.

Use the `no` parameter with this command to disable this function

Command Syntax

```
match ipv6 address prefix-list WORD
no match ipv6 address prefix-list (WORD|)
```

Parameter

`WORD`	IPv6 access list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match ipv6 address prefix-list mylist
```

match ipv6 next-hop

Use this command to specify the next-hop address to be matched. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid for BGP and IS-IS only.

Use the `no` parameter with this command to disable this function

Command Syntax

```
match ipv6 next-hop (X:X::X:X|WORD)
no match ipv6 next-hop (X:X::X:X|WORD|)
```

Parameters

`X:X::X:X`	IPv6 address of the next-hop.
`WORD`	IPv6 access list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match ipv6 next-hop 3ffe::1
```

match ipv6 next-hop prefix-list

Use this command to match entries of a prefix-list. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid for BGP and IS-IS only.

Use the `no` parameter with this command to disable this function

Command Syntax

```
match ipv6 next-hop prefix-list WORD
no match ipv6 next-hop prefix-list WORD
```

Parameters

`WORD`	IPv6 access list name.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match ipv6 next-hop prefix-list new
```

match ipv6 peer

Use this command to specify the match peer IPv6 address of a route.

Use the `no` parameter with this command to remove the specified match criterion.

Command Syntax

```
match ipv6 peer (<1-199>|<1300-2699>|WORD)
no match ipv6 peer (<1-199>|<1300-2699>|WORD|)
```

Parameter

`WORD`	IP access-list name.
`<1-199>`	IP access-list number (standard range).
`<1300-2699>`	IP access-list number (expanded range).

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#match ipv6 peer 123

(config-route-map)#no match ipv6 peer 123
```

match metric

Use this command to match a metric of a route. The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid for BGP, OSPF, RIP, and IS-IS only.

Use the `no` parameter with this command to disable this function

Command Syntax

```
match metric <0-4294967295>
no match metric (<0-4294967295>|)
```

Parameters

`<0-4261412864>` Metric value.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#match metric 888999
```

match origin

Use this command to match origin code. The origin attribute defines the origin of the path information. The `egp` parameter is indicated as an `e` in the routing table, and it indicates that the origin of the information is learned via EGP (Exterior Gateway Protocol). The `igp` parameter is indicated as an `i` in the routing table, and it indicates the origin of the path information is interior to the originating AS. The `incomplete` parameter is indicated as a `?` in the routing table, and indicates that the origin of the path information is unknown or learned through other means. If a static route is redistributed into BGP, the origin of the route is incomplete.

This command specifies the origin to be matched. If there is a match for the specified origin, and `permit` is specified when you created the route-map, the route is redistributed or controlled as specified by the set action. If the match criteria are met, and `deny` is specified, the route is not redistributed or controlled. If the match criteria are not met, the route is neither accepted nor forwarded, irrespective of `permit` or `deny` specifications.

The route specified by the policies might not be the same as specified by the routing protocols. Setting policies enable packets to take different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding specified in routing tables.

Note: This command is valid only for BGP.

Use the `no` parameter with this command to disable this matching.

Command Syntax

```
match origin (egp|igp|incomplete)
no match origin (egp|igp|incomplete|)
```

Parameters

`egp`	Remote exterior gateway protocol.
`igp`	Local internal gateway protocol.
`incomplete`	Unknown heritage.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map myroute deny 34
(config-route-map)#match origin egp
```

match route-type

Use this command to match an external route type. AS-external LSA is either Type-1 or Type-2. External type-1 matches only Type 1 external routes and external type-2 matches only Type 2 external routes.

Use the `no` parameter with this command to turn off the matching.

Command Syntax

```
match route-type external (type-1|type-2)
no match route-type external (type-1|type-2|)
```

Parameters

`type-1`	Match OSPF External Type 1 metric.
`type-2`	Match OSPF External Type 2 metric.

Default

By default, match route type external is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap1 permit 10
(config-route-map)#match route-type external type-1
```

match tag

Use this command to match the specified tag value.

Use the `no` parameter with this command to turn off the declaration.

Command Syntax

```
match tag <0-4294967295>
no match tag (<0-4294967295>|)
```

Parameters

`<0-4294967295>` Tag value.

Default

By default, match tag is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap1 permit 10
(config-route-map)#match tag 100
```

route-map

Use this command to enter route-map mode and to permit or deny match/set operations.

This command controls and modifies routing information to allow redistribution of routes. It has a list of `match` and `set` commands associated with it. The `match` commands specify the conditions under which redistribution is allowed, and the `set` commands specify the particular redistribution actions to be performed if the criteria enforced by match commands are met. Route maps are used for detailed control over route distribution between routing processes.

Route maps also allow policy routing, and might route packets to a different route than the obvious shortest path.

Use the `no` parameter with this command to turn off the declaration.

Command Syntax

```
route-map WORD (deny|permit) <1-65535>
no route-map WORD ((deny|permit) <1-65535>|)
```

Parameters

`WORD`	Route-map name (maximum size 63 characters).
`deny`	Route map deny set operations. If this parameter is specified, and the match criteria are met, the route is not redistributed, and any other route maps with the same map tag are not examined.
`permit`	Route map permit set operations. If this parameter is specified, and the match criteria are met, the route is redistributed as specified by the set actions. If the `match` criteria are not met, the next route map with the same tag is tested.
`<1-65535>`	Sequence to insert into or delete from an existing route-map.

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map route1 permit 1
(config-route-map)#
```

set aggregator

Use this command to set the AS number for the route map and router ID. An Autonomous System (AS) is a collection of networks under a common administration sharing a common routing strategy. It is subdivided by areas, and is assigned a unique 16-bit number. Use the `set aggregator` command to assign an AS number for the aggregator.

To use the `set aggregator` command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function

Command Syntax

```
set aggregator as <1-65535> A.B.C.D
no set aggregator as (<1-65535> A.B.C.D|)
```

Parameters

`<1-65535>`	AS number of aggregator.
`A.B.C.D`	IP address of aggregator.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#set aggregator as 43 10.10.0.3
```

set as-path

Use this command to modify an autonomous system path for a route. By specifying the length of the AS-Path, the router influences the best path selection by a neighbor. Use this command to prepend an AS path string to routes increasing the AS path length.

To use this command, users must first give the `match` and `set` commands configure the conditions for redistributing routes from one routing protocol to another:

- The `match` command specifies the match criteria under which redistribution is allowed for the current route-map.
- The `set` command specifies the set redistribution actions to be performed if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
set as-path prepend .<1-4294967295>
no set as-path prepend (.<1-4294967295>|)
```

Parameters

`<1-4294967295>` OcNOS prepends the 32 bit number to the AS path.

Default

None

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myroute permit 3
(config-route-map)#set as-path prepend 8 24
```

set atomic-aggregate

Use this command to set an atomic aggregate attribute.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function

Command Syntax

```
set atomic-aggregate
```

```
no set atomic-aggregate
```

Parameters

No default value is specified

Default

None

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set atomic-aggregate
```

set comm-list

Use this command to delete the matched communities from the community attribute of an inbound or outbound update when applying route-map.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
set comm-list (<1-99>|<100-199>|WORD) delete
no set comm-list (<1-99>|<100-199>|WORD) delete
```

Parameters

`<1-99>`	Standard community-list number.
`<100-199>`	Expanded community-list number.
`WORD`	Name of the community-list.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map myRM permit 3
(config-route-map)#set comm-list 34 delete
```

set community

Use this command to set the communities attribute. and group destinations in a certain community, as well as apply routing decisions according to those communities.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to delete the entry.

Command Syntax

```
set community [<65536-4294901759>|AA:NN|internet|local-AS|no-advertise|no-
  export|none] (additive|)
no set community [AA:NN|internet|local-AS|no-advertise|no-export|none] (additive|)
```

Parameters

`<65536-4294901759>`	Community number
`AA:NN`	The community number in aa:nn format.
`internet`	Internet.
`local-AS`	Do not send outside the local AS (well-known community).
`no-advertise`	Do not advertise this route to eBGP peers
`no-export`	Do not advertise this route to any peer.
`none`	Remove the community attribute from the prefixes that pass the route-map.
`additive`	Add to the existing community.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following examples show the use of the `set community` command with different parameters.

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set community no-export no-advertise

#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set community no-advertise
```

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set community 10:01 23:34 12:14 no-export
```

set dampening

Use this command to enable route-flap dampening and set parameters. Set the unreachability half-life time to be equal to, or greater than, reachability half-life time. The suppress-limit value must be greater than or equal to the reuse limit value.

Use the `no` parameter with this command to delete the entry.

Command Syntax

```
set dampening <1-45> <1-20000> <1-20000> <1-255> (<1-45>|)
no set dampening <1-45> <1-20000> <1-20000> <1-255> (<1-45>|)
```

Parameters

`<1-45>`	Reachability half-life time in minutes. The time for the penalty to decrease to one-half of its current value.
`<1-20000>`	Reuse-limit value. When the penalty for a suppressed route decays below the reuse value, the routes become unsuppressed.
`<1-20000>`	Suppress-limit value. When the penalty for a route exceeds the suppress value, the route is suppressed.
`<1-255>`	Max-suppress-time. Maximum time that a dampened route is suppressed.
`<1-45>`	Unreachability half-life time for penalty, in minutes.

Default

Default reachability half-life time: 15 minutes

Default reuse limit value: 750

Default suppress limit value: 2000

Default max-suppress value is 4 times the half-life time (60 minutes)

Default unreachability half-life time value: 15 minutes

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map R1 permit 24
(config-route-map)#set dampening 20 333 534 30
```

set extcommunity

Use this command to set an extended community attribute.

To use this command you must first have a match clause. The `match` and `set` commands define the conditions for redistributing routes from one routing protocol to another:

- The `match` command defines the match criteria under which redistribution is allowed for the current route-map.
- The `set` command defines the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
set extcommunity rt .AA:NN (additive|)
set extcommunity soo .AA:NN
set extcommunity cost (igp|pre-bestpath|) <0-255> <0-4294967295>
no set extcommunity rt (.AA:NN|) (additive|)
no set extcommunity soo (.AA:NN|)
no set extcommunity cost (igp|pre-bestpath|) <0-255> <0-4294967295>
```

Parameters

`rt`	Route target extended community.
`.AA:NN`	VPN extended community
`additive`	Add to the exsisting community.
`soo`	Site-of-origin extended community.
`cost`	Extended cost community.
`igp`	Compare following IGP cost comparison.
`pre-bestpath`	Compare following IGP cost comparison.
`<0-255>`	Community ID.
`<0-4294967295>`	Cost.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
```

```
(config)#route-map rmap1 permit 3
(config-route-map)#set extcommunity rt 06:01

#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set extcommunity rt 0.0.0.6:01

#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set extcommunity soo 06:01

#configure terminal
(config-route-map)#route-map rmap1 permit 3
(config-route-map)#set extcommunity soo 0.0.0.6:01
```

set interface null0

Use this command to discard routes based on policy/rules configured for a route map.

Route maps can be applied to BGP neighbors. When this command is given for a route map and that route map is applied to a BGP neighbor, the discard route entries are added by BGP for the prefix permitted by the route map.

Use the `no` form of this command to not discard routes based on policy/rules configured for a route map.

Command Syntax

```
set interface null0
no set interface null0
```

Parameters

None

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

To discard 1.1.1.1/24 from BGP neighbor 30.1.1.1:

```
#configure terminal
(config)#ip prefix-list myPrefixList seq 5 permit 1.1.1.1/24 eq 24
(config)#route-map myRM permit 1
(config-route-map)#match ip address prefix-list myPrefixList
(config-route-map)#set interface null0
(config-route-map)#exit
(config)#router bgp 200
(config-router)#neighbor 30.1.1.1 remote-as 200
(config-router)#neighbor 30.1.1.1 route-map myRM in
```

© 2025 IP Infusion Inc. Proprietary

set ip next-hop

Use this command to set the specified next-hop value.

Use the `no` parameter with this command to turn off the setting.

Note: This command is valid for BGP, OSPF, and RIP only.

Command Syntax

```
set ip next-hop A.B.C.D
no set ip next-hop (A.B.C.D|)
```

Parameter

`A.B.C.D`	IP address of the next-hop.

Default

By default, set ip next hop A.B.C.D is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#set ip next-hop 10.10.0.67
```

set ipv6 next-hop

Use this command to set a next hop-address.

Use the `no` parameter with this command to delete an entry.

Note: This command is valid for BGP and OSPFv3 only.

Command Syntax

```
set ipv6 next-hop X:X::X:X
set ipv6 next-hop local X:X::X:X
no set ipv6 next-hop (X:X::X:X|)
no set ipv6 next-hop local (X:X::X:X|)
```

Parameters

`X:X::X:X`	Global IPv6 address of the nexthop.
`local`	Local IPv6 address of the nexthop.

Default

By default, set ipv6 next hop X:X::X:X is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set ipv6 next-hop local fe80::203:47ff:fe97:66dc
```

set level

Use this command to set the IS-IS level to export a route.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
set level (level-1|level-2|level-1-2)
no set level (level-1|level-2|level-1-2|)
```

Parameters

`level-1`	Export into a level-1 area.
`level-2`	Export into a level-2 sub-domain.
`level-1-2`	Export into level-1 and level-2.

Default

By default, set level is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set level level-1
```

set local-preference

Use this command to set the BGP local preference path attribute.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
set local-preference <0-4294967295>
no set local-preference (<0-4294967295>|)
```

Parameters

`<0-4294967295>` Tag value for destination routing protocol.

Default

By default, set local preference is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set local-preference 12
```

set metric

Use this command to set a metric value for a route and influence external neighbors about the preferred path into an Autonomous System (AS). The preferred path is the one with a lower metric value. A router compares metrics for paths from neighbors in the same ASs. To compare metrics from neighbors coming from different ASs, use the `bgp always-compare-med` command.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
set metric WORD
no set metric (WORD|)
```

Parameters

`WORD`	Metric value [+/-] <1-4294967295>

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set metric 600
```

set metric-type

Use this command to set the metric type for the destination routing protocol. Select a type to be either Type-1 or Type-2 in the AS-external-LSA when the route-map matches the condition.

Note: This command is for OSPF, OSPFv3, or IS-IS only.

Use the `no` parameter with this command to return to the default.

Command Syntax

```
set metric-type (internal|external)
set metric-type (type-1|type-2)
no set metric-type (internal|external|)
no set metric-type (type-1|type-2|)
```

Parameters

`external`	IS-IS external metric type.
`internal`	IS-IS internal metric type.
`type-1`	OSPF external type 1 metric.
`type-2`	OSPF external type 2 metric

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

In this example the metric type of the destination protocol is set to OSPF external Type 1.

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set metric-type type-1
```

set origin

Use this command to set the BGP origin code. The origin attribute defines the origin of the path information.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to delete an entry.

Command Syntax

```
set origin (egp|igp|incomplete)
no set origin (egp|igp|incomplete|)
```

Parameters

`egp`	Learned through an Exterior Gateway Protocol.
`igp`	Interior to the originating AS. This happens when an Internal Gateway Protocol is redistributed into BGP.
`incomplete`	Unknown or learned through some other means. This happens when a static route is redistributed in BGP and the origin of the route is incomplete.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set origin egp
```

set originator-id

Use this command to set the originator ID attribute.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

If the packets do not match any of the defined criteria, they are routed through the normal routing process.

Use the `no` parameter with this command to disable this function

Command Syntax

```
set originator-id A.B.C.D
no set originator-id (A.B.C.D|)
```

Parameter

`A.B.C.D`	IP address of originator.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set originator-id 1.1.1.1
```

set tag

Use this command to set a tag value. The parameter is the route tag that is labeled by another routing protocol (BGP or other IGP when redistributing), because AS-external-LSA has a route-tag field in its LSAs. In addition, when using route-map, OcNOS can tag the LSAs with the appropriate tag value. Sometimes the tag matches with using route-map, and sometimes, the value may be used by another application.

Use the `no` parameter with this command to unset a tag value.

Command Syntax

```
set tag <0-4294967295>
no set tag (<0-4294967295>|)
```

Parameter

`<0-4294967295>` Tag value for destination routing protocol.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

In the following example the tag value of the destination routing protocol is set to 6:

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set tag 6
```

set weight

Use this command to set weights for the routing table.

The weight value is used to assist in best path selection. It is assigned locally to a router. When there are several routes with a common destination, the routes with a higher weight value are preferred.

To use this command, you must first have a match clause. `Match` and `set` commands set the conditions for redistributing routes from one routing protocol to another. The `match` command specifies the match criteria under which redistribution is allowed for the current route-map. The `set` command specifies the set redistribution actions to be performed, if the match criteria are met.

Note: This command is valid only for BGP.

Use the `no` parameter with this command to delete an entry.

Command Syntax

```
set weight <0-4294967295>
no set weight (<0-4294967295>|)
```

Parameter

`<0-4294967295>` Weight value.

Default

No default value is specified

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following configuration, all routes that apply to access-list 10 will have the weight set at 400. If the packets do not match any of the defined criteria, they are routed through the normal routing process.

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#match as-path 10
(config-route-map)#set weight 400
```

show route-map

Use this command to display route-map information.

Command Syntax

```
show route-map (|WORD)
```

Parameters

WORD	Route-map name (maximum size 63 characters)

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of the `show route-map` command.

```
#show route-map
route-map myRM, permit, sequence 1
  Match clauses:
    metric 200
  Set clauses:
    metric 60
#
```

Table 11 explains the output fields.

Table 11: show route-map details

Field	Description
route-map	Name of a route map.
permit	Routes that match the match clauses are redistributed according to the conditions defined by the set clauses.
sequence	Position of this route map in sequence of route-maps with the same name.
Match clauses	Routes that match the conditions defined by the match clause are redistributed according to the conditions defined by the set clauses.
Set clauses	Routes that pass the match clause are redistributed according to the conditions defined by the set clauses.

show running-config route-map

Use this command to display the running system status and configuration details for route-maps.

Command Syntax

```
show running-config route-map
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#show running-config route-map
!
route-map abc deny 2
 match community 2
!
route-map abc permit 3
 match route-type external type-2
 set metric-type type-1
!
```

CHAPTER 3 Unidirectional Link Detection Commands

This section describes the Unidirectional Link Detection (UDLD) commands.

- udld
- udld message-time
- udld mode
- udld state
- show udld
- show udld interface

udld

Use this command to enable or disable the UDLD feature globally.

Command Syntax

```
udld  (enable | disable)
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS Version 5.0

Examples

```
(config)#udld enable
```

udld message-time

Use this command to set the UDLD message interval.

Command Syntax

```
udld message-time <7-90>
```

Parameters

<7-90>	Interval time in seconds

Default

15 seconds

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS Version 5.0.

Examples

```
config)#udld message-time 50
```

udld mode

Use this command to configure UDLD mode as aggressive or normal.

Command Syntax

```
udld mode (aggressive | normal)
```

Parameters

`aggressive`	Aggressive mode
`normal`	Normal mode

Default

N/A

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS Version 5.0.

Examples

```
(config-if)#udld mode aggressive
```

udld state

Use this command to enable or disable the UDLD feature for an interface.

Command Syntax

```
udld state (enable | disable)
```

Parameters

None

Default

Disabled

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS Version 5.0.

Examples

```
(config)#int xe7
(config-if)#udld state enable
```

show udld

Use this command to display UDLD statistic for all interface.

Command Syntax

```
show udld
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS-SP version 5.0.

Examples

```
#show udld
UDLD                    : Enable
Message Interval(sec)   : 15
Port    UDLD Status     Mode            Link-Status
------------------------------------------------
xe7     Enable          Normal          Bi-Directional
```

Table 3-1 explains the output fields.

Table 3-1: show udld output fields

Field	Description
UDLD	Whether UDLD is enabled or disabled
Message Interval	Message interval in seconds
Port	Interface name
UDLD Status	Whether UDLD is enabled or disabled on the interface
Mode	Whether the mode is aggressive or normal
Link-Status	State of the link: Unknown Loop-Back Neighbor Mismatch Unidirectional Undetermined Bi-Directional

show udld interface

Use this command to display UDLD settings for particular interface.

Command Syntax

```
show udld interface IFNAME
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS-SP version 5.0.

Examples

```
#show udld interface xe14
UDLD Status           : Enable
UDLD Mode             : Aggressive
Link-State            : Bi-Directional
#
```

Table 3-2 explains the output fields.

Table 3-2: show udld interface output fields

Field	Description
UDLD Status	Whether UDLD is enabled or disabled
UDLD Mode	Whether the mode is aggressive or normal
Link-State	State of the link: Unknown Loop-Back Neighbor Mismatch Unidirectional Undetermined Bi-Directional

CHAPTER 4 Object Tracking Commands

This chapter describes the Object Tracking commands:

- track ip sla reachability
- delay up down
- object-tracking
- show track
- show track <1-500>
- show track summary
- show running-config track

track ip sla reachability

Use this command to configure an Object for tracking using IP SLA.

Use the `no` form of this command to delete to object tracking

Command Syntax

```
track <1-500> ip sla <1-65535> reachability)
no track <1-500> ip sla <1-65535> reachability
```

Parameters

`object-number (1-500)` Identifier for the tracked object

`ip-sla-number (1-65535)` Ientifier for IP SLA association with tracking object

Command Mode

Configuration mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
#configure terminal
OcNOS(config)#track 1 ip sla 1 reachability
OcNOS(config-object-track)#commit

OcNOS(config)#no track 1
OcNOS(config)#commit
```

delay up down

Use This command is used to delay the state change notification of Object tracking.

Use the `no` form of this command to remove delay the state change notification of Object

Command Syntax

```
delay (up <1-9999>|)(down <1-9999>|)
no delay (|up|down)
```

Parameters

`<1-999>`	Delay in Notification in seconds.

Default

NA

Command Mode

Object tracking Mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS(config-object-track)#delay up 10 down 20
OcNOS(config-object-track)#no delay
OcNOS(config-object-track)#commit
OcNOS(config-object-track)#
OcNOS(config-object-track)#delay down 10
OcNOS(config-object-track)#commit
OcNOS(config-object-track)#no delay down
OcNOS(config-object-track)#commit
OcNOS(config-object-track)#
OcNOS(config-object-track)#delay up 10
OcNOS(config-object-track)#commit
OcNOS(config-object-track)#no delay up
OcNOS(config-object-track)#commit
OcNOS(config-object-track)#
```

show track

Use this command to display Sham link information.

Command Syntax

```
show track
```

Parameters

None

Default

NA

Command Mode

Exec mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#sh track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is DOWN
    0 changes, last change : 2021 Dec 11 05:20:23
OcNOS#
```

show track <1-500>

Use this command to display Sham link information.

Command Syntax

```
show track <1-500>
```

Parameters

<1-500>	object identifier

Default

NA

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#sh track 2
TRACK Id: 2
  IP SLA 2 reachability
  Reachability is DOWN
    0 changes, last change : 2021 Dec 11 05:29:49
OcNOS#
```

show track summary

Use this command to display the summary of all object tracking.

Command Syntax

```
show track summary
```

Parameters

NA

Default

NA

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#sh track summary
Object Tracking Summary
ID       Type         Type-Identifier       State
-----------------------------------------------------
1      ip-sla             1                 DOWN
2      ip-sla             2                 DOWN
OcNOS#
```

show running-config track

Use this command to display object tracking running configuration alone.

Command Syntax

```
show running-config track
```

Parameters

NA

Default

NA

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#sh running-config track
track 1 ip sla 1 reachability
 delay up 20
!
track 2 ip sla 2 reachability
!
OcNOS#
```

CHAPTER 5 IP Service Level Agreements Commands

IP Service Level Agreements (SLAs) is a diagnostic method which generates and analyses the traffic between an OcNOS device and your network. IP SLA monitors and reports network performance data which helps you to identify the actual root cause of a problem when the performance level drops.

This chapter describes the commands used to manage the IP SLA for ICMP echo.

- clear ip sla statistics
- frequency
- icmp-echo
- ip sla
- ip sla schedule
- show ip sla statistics
- show ip sla summary
- show running-config ip sla
- threshold
- timeout

© 2025 IP Infusion Inc. Proprietary

clear ip sla statistics

Use this command to clear the IP SLA statistics.

Command Syntax

```
clear ip sla statistics <1-65535>
```

Parameters

1-65535	IP SLA identifier

Default

N/A

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#clear ip sla statistics 1
```

frequency

Use this command to configure the frequency/interval to send ICMP echo packets one by one.

Use the `no` form of this command to remove the configured ICMP echo frequency.

Command Syntax

```
frequency <1-60>
no frequency
```

Parameters

`1-60`	Frequency in seconds

Default

5 seconds

Command Mode

IP SLA ICMP Echo mode (config-ip-sla-echo)

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 10.12.28.1 source-interface xe1
(config-ip-sla-echo)#frequency 3
```

icmp-echo

Use this command to select and configure the ICMP echo SLA operation. ICMP echo packets are constructed in the device and sent to the destination address that you specify. These packets are transferred on a specific interface by setting the `source-interface` parameter.

Use the `no` form of this command to un-configure or remove the configured ICMP echo measurement sessions.

Command Syntax

```
icmp-echo (ipv4 A.B.C.D|ipv6 X:X::X:X|HOSTNAME) (source-interface IFNAME|)
no icmp-echo (ipv4 A.B.C.D | ipv6 X:X::X:X | HOSTNAME)
```

Parameters

`A.B.C.D`	IPv4 address
`X:X::X:X`	IPv6 address
`HOSTNAME`	Host name
`IFNAME`	Source interface name

Default

N/A

Command Mode

IP SLA mode (config-ip-sla)

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 10.12.28.1 source-interface xe1
(config-ip-sla-echo)#
```

ip sla

Use this command to create an IP SLA instance. One instance maps to a single SLA operation. You can create multiple SLA operations to perform multiple similar or different SLA operations.

Use the `no` form of this command to remove a configured IP SLA configurations.

Command Syntax

```
ip sla <1-65535>
no ip sla <1-65535>
```

Parameters

`1-65535`	IP SLA identifier

Default

N/A

Command Mode

Configuration mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#
```

ip sla schedule

Use this command to schedule an IP SLA operation by associating a time-range object with the IP SLA operation.

Use the `no` form of this command to stop the configured IP SLA measurement.

Command Syntax

```
ip sla schedule <1-65535> time-range WORD (vrf (NAME)|)
```

Parameters

`<1-65535>`	IP SLA identifier.
`time-range`	Time Range
`TR_NAME`	Time range name that you set with the time-range command.
`vrf`	VPN Routing/Forwarding instance
`NAME`	VPN Routing/Forwarding instance name. Maximum limit 32 characters

Default

N/A

Command Mode

Configuration mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#ip sla schedule 1 time-range t1 vrf v1
```

show ip sla statistics

Use this command to display the statistics of IP SLA measurement.

Command Syntax

```
show ip sla statistics (1-65535) detail
```

Parameters

1-65535	IP SLA identifier.

Default

N/A

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show ip sla statistics 1 detail
========================================
            IP SLA Statistics
========================================
  IP SLA ID                : 1
  Start Time               : 2021 Aug 30 17:40:04
  Elapsed time(milli sec)  : 46015
  Packets Sent             : 23
  Packets Received         : 23
  Packet Loss(%)           : 0.0000
  Invalid Tests            : 0
  Round Trip Delay(usec)
      Minimum              : 1000
      Maximum              : 1000
      Average              : 1000
```

Table 5-3 explains the output fields.

Table 5-3: show ip sla statistics fields

Field	Description
IP SLA ID	IP SLA Identifier (1-65535)
Start Time	Measurement start time
Elapsed time(milli sec)	Time taken to complete the measurement in milliseconds
Packets Sent	Number of packet sent

Table 5-3: show ip sla statistics fields (Continued)

Field	Description
Packets Received	Number of packet received
Packet Loss(%)	Packet lost in percentage
Invalid Tests	Received ICMP echo reply packets after configured threshold limit will be marked as invalid tests
Round Trip Delay(usec)	Round trip delay between ICMP echo request and ICMP echo reply: minimum, maximum and average round trip delay in microseconds

show ip sla summary

Use this command to display the summary of all IP SLA measurements.

Command Syntax

```
show ip sla summary
```

Parameters

None

Default

N/A

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show ip sla summary
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                       (usec)        Code        Run
---------------------------------------------------------------------
^1      icmp-echo    20.2.2.3           0           OK        2021 Aug 23 13:53:37
```

Table 5-4 explains the output fields.

Table 5-4: show ip sla summary fields

Field	Description
ID	IP SLA Identifier (1-65535)
Type	Measurement type
Destination	Destination address
Stats (usec)	Round trip time in microseconds for the measurement
Return Code	Measurement status
Last Run	Measurement last run date and time

show running-config ip sla

Use this command to display the IP SLA running configuration alone.

Command Syntax

```
show running-config ip sla
```

Parameters

None

Default

N/A

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show running-config ip sla
ip sla 1
 icmp-echo ipv4 20.2.2.3
  frequency 2
  threshold 2000
  timeout 5000
ip sla schedule 1 time-range t1 vrf v1
```

threshold

Use this command to configure the threshold for every ICMP echo packet.

Use the `no` form of this command to remove the configured ICMP echo threshold.

Command Syntax

```
threshold <1000-60000>
no threshold
```

Parameters

`1000-60000`	Threshold in milliseconds.

Default

10000 milliseconds

Command Mode

IP SLA ICMP Echo mode (config-ip-sla-echo)

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 10.12.28.1 source-interface xe1
(config-ip-sla-echo)#threshold 5000
```

timeout

Use this command to configure the timeout for every ICMP echo packet. Any packet arriving beyond this interval is considered to be lost.

Use the `no` form of this command to remove the configured ICMP echo timeout.

Command Syntax

```
timeout <1000-60000>
no timeout
```

Parameters

`1000-60000`	Timeout in milliseconds.

Default

10000 milliseconds

Command Mode

IP SLA ICMP Echo mode (config-ip-sla-echo)

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 10.12.28.1 source-interface xe1
(config-ip-sla-echo)#timeout 5000
```

CHAPTER 6 Source Interface Commands

This chapter is a reference for source interface commands. The source Interface feature routes management traffic to a dedicated interface using `iptables` NAT rules.

The source interface feature is supported for the protocols shown in Table 6-5.

Table 6-5: Source interface protocols and port numbers

Protocol	Default port number
Tacacs+	49
Radius	1812 and 1813
Snmp	161 and 162
Ntp	123
Syslog	514

Note: Because management applications are allowed only on the default and management VRF, the commands in this chapter are supported on the "management" and "default" VRFs only.

This chapter contains these commands:

- ip source-interface
- ipv6 source-interface
- show ip source-interface detail
- show ipv6 source-interface detail
- show running-config ip source-interface
- show running-config ipv6 source-interface

ip source-interface

Use this command to configure the IPv4 source interface for a protocol.

Use the `no` form of this command to remove the IPv4 source interface for a protocol.

Note: It is possible that the router may establish an outgoing TCP connection using an interface that does not have a valid or routable IP address. In such case, the user must specify the address of a different interface to use as the source IP address for the outgoing connection. For this scenario, the command ip source-interface or ipv6 source-interfaceis used.

Command Syntax

```
ip source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius) (port (0|<1025-
  65535>)|)(vrf (NAME|management)|)
no ip source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius)
```

Parameters

`IFNAME`	Interface name (lo or physical interface)
`tacacs+`	Terminal Access Controller Access Control System
`ntp`	Network Time Protocol
`snmp`	Simple Network Management Protocol
`syslog`	syslog
`radius`	Remote Authentication Dial-In User Service
`<1025-65535>`	Port number. Default value is as per the protocol.
`vrf management`	Defines the management VRF instance.
`vrf NAME`	Specify the user-defined VRF instance name.

Default

The default port 0 is assigned to the protocol.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0. Added VRF NAME parameter in OcNOS version 6.6.0

Example

```
#configure terminal
(config)#ip source-interface xe1 ntp
(config)#ip source-interface xe2 radius port 1025
(config)#ip source-interface xe3 syslog port 65535 vrf management
```

ipv6 source-interface

Use this command to configure the IPv6 source interface for a protocol.

Use the `no` form of this command to remove the IPv6 source interface for a protocol.

Command Syntax

```
ipv6 source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius) (port (0|<1025-
  65535>)|) (vrf (NAME|management)|)
no ipv6 source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius)
```

Parameters

`IFNAME`	Interface name (lo or physical interface)
`tacacs+`	Terminal Access Controller Access Control System protocol
`ntp`	Network Time Protocol
`snmp`	Simple Network Management Protocol
`syslog`	syslog
`radius`	Remote Authentication Dial-In User Service
`<1025-65535>`	Port number. Default value is as per the protocol.
`vrf management`	Defines the management VRF instance.
`vrf NAME`	Specify the user-defined VRF instance name.

Default

The default port 0 is assigned to the protocol.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)#ipv6 source-interface xe1 ntp
(config)#ipv6 source-interface xe2 radius port 1025
(config)#ipv6 source-interface xe3 syslog port 65535 vrf management
```

© 2025 IP Infusion Inc. Proprietary

show ip source-interface detail

Use this command to display the IPv4 source interface status in detail.

Command Syntax

```
show ip source-interface detail
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#show ip source-interface detail
Source-Interface Detailed Information
=====================================
  Protocol : tacacs+
  Interface : lo
  Address : 1.1.1.1
  Status : Active
  VRF Name : Default

  Protocol : radius
  Interface : lo
  Address : 1.1.1.1
  Status : Active
  VRF Name : Default
```

Table 6-6 explains the output fields.

Table 6-6: Output fields

Field	Description
Protocol	tacacs+, ntp, snmp, syslog, or radius
Interface	Interface name (lo or physical interface)
Address	IP address
Status	Whether active or inactive
VRF Name	Virtual Routing and Forwarding name

show ipv6 source-interface detail

Use this command to display the IPv6 source interface status in detail.

Command Syntax

```
show ipv6 source-interface detail
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#show ipv6 source-interface detail
Source-Interface Detailed Information
=====================================
  Protocol : tacacs+
  Interface : lo
  Address : ::1
  Status : Active
  VRF Name : Default

  Protocol : radius
  Interface : lo
  Address : ::1
  Status : Active
  VRF Name : Default
```

Table 6-6 explains the output fields.

show running-config ip source-interface

Use this command to display the IPv4 source interface running configuration.

Command Syntax

```
show running-config ip source-interface
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 4.0

Example

```
#show running-config ip source-interface
ip source-interface lo tacacs+ port 1025
ip source-interface lo radius
ip source-interface lo.management ntp vrf management
ip source-interface lo.management syslog port 1026 vrf management
ip source-interface ge3 snmp
```

show running-config ipv6 source-interface

Use this command to display the IPv6 source interface running configuration.

Command Syntax

```
show running-config ipv6 source-interface
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#show running-config ipv6 source-interface
ipv6 source-interface lo tacacs+ port 1025
ipv6 source-interface lo radius
ipv6 source-interface lo.management ntp vrf management
ipv6 source-interface lo.management syslog port 1026 vrf management
ipv6 source-interface ge3 snmp
```

Border Gateway Protocol Configuration

CHAPTER 1 BGP

This chapter contains basic Border Gateway Protocol configuration examples.

Enable BGP Routers in the Same Autonomous System

Figure 1-9 shows the minimum configuration required to enable BGP on an interface. R1 and R2 are two routers belonging to the same AS, AS200, connecting to network 10.10.10.0/24. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

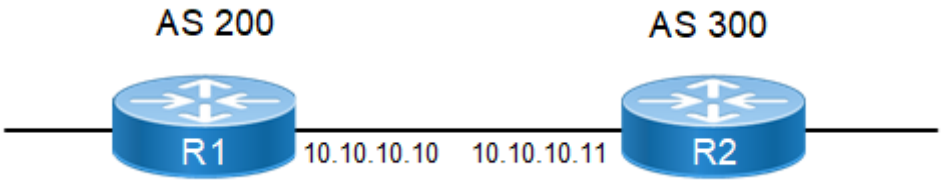

Figure 1-9: Routers in the Same Autonomous System

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of R1.
`(config-router)#neighbor 10.10.10.11 remote-as 200`	Define BGP neighbors, and establish a TCP session. `10.10.10.11` is the IP address of the neighbor (R2), and `200` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.11 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)# exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of `R2`.
`(config-router)#neighbor 10.10.10.10 remote-as 200`	Define BGP neighbors, and establish a TCP session. `10.10.10.10` is the IP address of the neighbor (`R1`), and `200` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.10 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)# exit-address-family`	Exit address-family mode.

Validation

```
#show ip bgp summary
BGP router identifier 192.168.52.2, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.10.10.11              4   200  387        390       1      0      0
00:00:04              0

Total number of neighbors 1

Total number of Established sessions 1

#show ip bgp neighbors
BGP neighbor is 10.10.10.11, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  BGP state = Established, up for 00:01:41
  Last read 00:00:11, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 5 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.10, Local port: 179
Foreign host: 10.10.10.11, Foreign port: 33931
Nexthop: 10.10.10.10
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Enable BGP Between Different Autonomous Systems

This example shows the minimum configuration required for enabling BGP on an interface, when the routers belong to different autonomous systems. R1 and R2 are two routers in different autonomous system, AS200 and AS300, connecting to network 10.10.10.0/24.

Topology

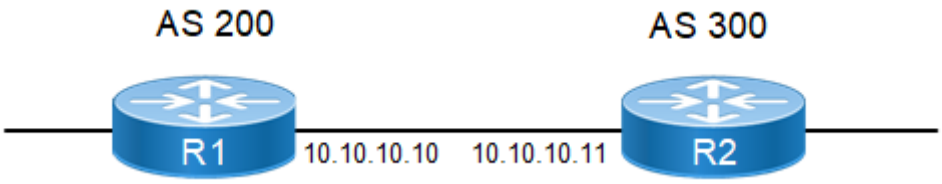

Figure 1-10: Routers in Different Autonomous Systems

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of R1.
`(config-router)#neighbor 10.10.10.11 remote-as 300`	Define BGP neighbors, and establish a TCP session. `10.10.10.11` is the IP address of the neighbor (R2), and `300` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.11 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)# exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 300`	Define the routing process. The number `300` specifies the AS number of `R2`.
`(config-router)#neighbor 10.10.10.10 remote-as 200`	Define BGP neighbors, and establish a TCP session. `10.10.10.10` is the IP address of the neighbor (R1), and `200` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.10 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)# exit-address-family`	Exit address-family mode.

Validation

```
#show ip bgp neighbors
BGP neighbor is 10.10.10.10, remote AS 200, local AS 300, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  BGP state = Established, up for 00:00:15
  Last read 00:00:15, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 2 messages, 0 notifications, 0 in queue
  Sent 2 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
```

© 2025 IP Infusion Inc. Proprietary

```
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.11, Local port: 56091
Foreign host: 10.10.10.10, Foreign port: 179
Nexthop: 10.10.10.11
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip bgp summary
BGP router identifier 192.168.52.3, local AS number 300
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS    MsgRcv    MsgSen TblVer    InQ    OutQ    Up/
Down    State/PfxRcd
10.10.10.10              4    200    3          3       1       0       0
00:00:50              0

Total number of neighbors 1

Total number of Established sessions 1
```

Route-Map

Use route maps to filter incoming updates from a BGP peer. In this example, the prefix-list `ABC` on R1 is configured to deny entry of any routes with the IP address 1.1.1.0/M (M = 26, 27, 28). To test the filter, R2 is configured to generate network addresses 1.1.1.0/27 and 1.1.2.0/24. To verify, use the `show ip bgp` command on R1; it displays R1 receiving updates from only 1.1.2.0/24.

Topology

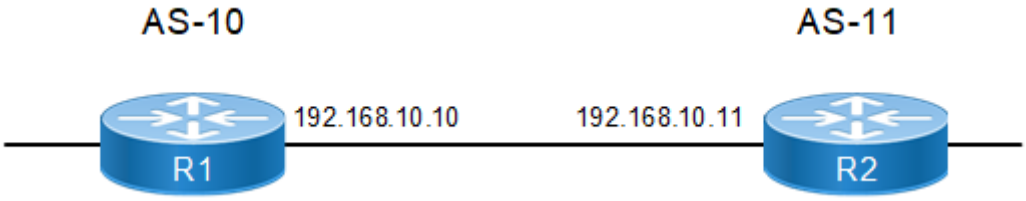

Figure 1-11: Configure Route-Map

R1

`#configure terminal`	Enter configure mode.
`(config)#ip prefix-list ABC`	Create an entry in the prefix-list. The ABC parameter is the name of the map that is created above. 5 specifies the sequence number or position of this specific route map. deny specifies the packets are to be rejected. 26 and 28 are the minimum and maximum prefix lengths to be matched.
`(config-ip-prefix-list)#seq 5 deny 1.1.1.0/24 ge 26 le 28`	5 specifies the sequence number or position of this specific route map. deny specifies the packets are to be rejected. 26 and 28 are the minimum and maximum prefix lengths to be matched.
`(config-ip-prefix-list)#seq 10 permit any`	10 specifies the sequence number or position of this specific route map. The permit parameter any specifies accept all packets of any length.
`(config-ip-prefix-list)#exit`	Exit the prefix-list mode
`(config)#route-map ABC permit 1`	Enter Route-map mode to set the match operation.
`(config-route-map)#match ip address prefix-list ABC`	Set the match criteria. In this case, if the route-map name matches ABC, the packets from the first sequence are denied.
`(config-route-map)#exit`	Exit Route-map mode, and return to Configure mode.
`(config)#router bgp 10`	Define the routing process, and establish a TCP session. The number 10 specifies the AS number of R1.
`(config-router)#neighbor 192.168.10.11 remote-as 11`	Define BGP neighbors, and establish a TCP session. 192.168.10.11 is the IP address of the neighbor (R2), and 11 is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 192.168.10.11 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#neighbor 192.168.10.11 route-map ABC in`	Apply a route map to routes. 192.168.10.11 specifies the IP address of BGP neighbor. The ABC parameter is the name of the route map, and in specifies that the access list applies to incoming advertisements.
`(config-router-af)#exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 1.1.1.1/27 secondary`	Specify the interface address.
`(config-if)#ip address 1.1.2.1/24 secondary`	Specify the interface address.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#router bgp 11`	Define the routing process, and establish a TCP session. The number `11` specifies the AS number of R2.
`(config-router)#neighbor 192.168.10.10 remote-as 10`	Define BGP neighbors, and establish a TCP session. `192.168.10.10` is the IP address of the neighbor (R1), and `10` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode

`(config-router-af)# neighbor 192.168.10.10 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#network 1.1.1.0/27`	Specify the network to be advertised by the BGP routing process.
`(config-router-af)#network 1.1.2.0/24`	Specify the network to be advertised by the BGP routing process.
`(config-router-af)#exit-address-family`	Exit router mode.

Validation

```
#show  ip bgp
BGP table version is 2, local router ID is 192.168.52.2
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   1.1.2.0/24       192.168.10.11         0         100            0       11
i

Total number of prefixes 1
```

Route Reflector

The configurations in this section apply to BGP Route Reflectors (RR).

Reduce the iBGP Mesh Inside an Autonomous System

Use Route Reflectors to reduce the iBGP mesh inside an Autonomous System (AS).

Topology

In this example, R2, R5, and R4 would have to maintain a full mesh among themselves, but by making R5 the Route Reflector, R2 (Client1) has an iBGP session with the RR only, but not with R4 (Client 2). The routes learned from R2 are advertised to the other clients, and to iBGP peers outside the cluster; the iBGP routes learned from iBGP peers outside the cluster are advertised to R2. This reduces the iBGP peer connections in AS1.

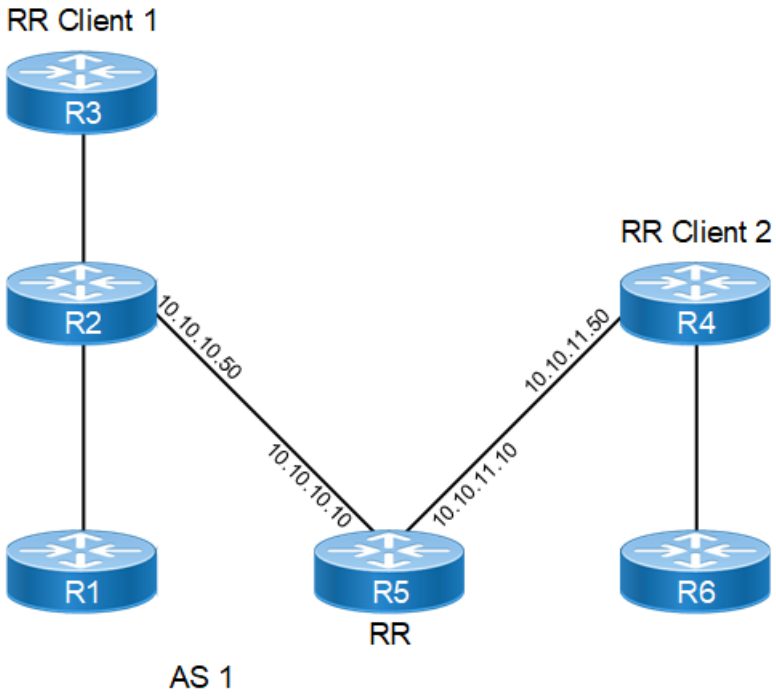

Figure 1-12: BGP Route Reflector

RR (R5)

`#configure terminal`	Enter configure mode.
`(config)#router bgp 1`	Define the routing process. The number 1 identifies the AS number of R5.
`(config-router)#neighbor 10.10.10.50 remote-as 1`	Define the BGP neighbor, and establish a TCP session. 10.10.10.50 is the IP address of one of the neighbors (R2), and 1 is the neighbor's AS number.
`(config-router)#neighbor 10.10.11.50 remote-as 1`	Define the BGP neighbor, and establish a TCP session. 10.10.11.50 is the IP address of one of the neighbors (R4), and 1 is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.50 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#neighbor 10.10.11.50 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#neighbor 10.10.10.50 route-reflector-client`	Configure R5 as the Route-Reflector (RR) and neighbor R2 as its client.
`(config-router-af)#neighbor 10.10.11.50 route-reflector-client`	Configure R5 as the Route-Reflector (RR) and neighbor R4 as its client.
`(config-router-af)#exit-address-family`	Exit address-family mode.

© 2025 IP Infusion Inc. Proprietary

RR Client 1 (R2)

`(config)#router bgp 1`	Define the routing process. The number `1` specifies the AS number of R2.
`(config-router)#neighbor 10.10.10.10 remote-as 1`	Define the BGP neighbor, and establish a TCP session. `10.10.10.10` is the IP address of the neighbor (R5), and `1` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.10.10 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#exit`	Exit router mode.

RR Client 2 (R4)

`(config)#router bgp 1`	Define the routing process. The number `1` identifies the AS number of `R4`.
`(config-router)#neighbor 10.10.11.10 remote-as 1`	Define BGP neighbor, and establish a TCP session. `10.10.11.10` is the IP address of the neighbor (R5), and `1` is the neighbor's AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.10.11.10 activate`	Activate the neighbor in the Ipv4 address family.
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#exit`	Exit router mode.

Validation

R5

```
#show ip bgp neighbors
BGP neighbor is 10.10.10.50, remote AS 1, local AS 1, internal link
  BGP version 4, local router ID 192.160.50.3, remote router ID 10.12.4.152
  BGP state = Established, up for 00:01:04
  Last read 00:01:04, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
```

```
Local host: 10.10.10.10, Local port: 47983
Foreign host: 10.10.10.50, Foreign port: 179
Nexthop: 10.10.10.10
Nexthop global: fe80::a00:27ff:fe09:fd25
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 10.10.11.50, remote AS 1, local AS 1, internal link
  BGP version 4, remote router ID 10.12.4.197
  local router ID 192.160.50.3
  BGP state = Established, up for 00:01:04
  Last read 00:01:04, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.11.10, Local port: 39851
Foreign host: 10.10.11.50, Foreign port: 179
Nexthop: 10.10.11.10
Nexthop global: fe80::a00:27ff:fe52:45f6
Nexthop local: ::
BGP connection: non shared network
```

R3

```
#show ip bgp neighbors
BGP neighbor is 10.10.11.10, remote AS 1, local AS 1, internal link
  BGP version 4, local router ID 192.160.50.4, remote router ID 10.12.4.185
  BGP state = Established, up for 00:00:56
  Last read 00:00:56, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 3 messages, 0 notifications, 0 in queue
  Sent 3 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
```

```
Local host: 10.10.11.50, Local port: 179
Foreign host: 10.10.11.10, Foreign port: 39851
Nexthop: 10.10.11.50
Nexthop global: fe80::a00:27ff:fe42:fb7a
Nexthop local: ::
BGP connection: non shared network
```

R2

```
#show ip bgp neighbors
BGP neighbor is 10.10.10.10, remote AS 1, local AS 1, internal link
  BGP version 4, local router ID 192.160.50.2, remote router ID 10.12.4.185
  BGP state = Established, up for 00:01:23
  Last read 00:01:23, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.50, Local port: 179
Foreign host: 10.10.10.10, Foreign port: 47983
Nexthop: 10.10.10.50
Nexthop global: fe80::a00:27ff:fe9c:f35d
Nexthop local: ::
BGP connection: non shared network
```

Multiple Route Reflectors

The basic rule of BGP is that a BGP speaker cannot advertise a route to an iBGP neighbor if that route was learned from another iBGP neighbor. Configuring a route reflector provides a means to circumvent this rule. The entire route reflector process is transparent to the clients, and no configuration is necessary on these clients.

Whenever an iBGP-speaking router receives a route update, it forwards the route to the neighbor without changing the nexthop IP address, thus making it an unreachable route, unless verified by an iGP (for example, neighbor x.x.x.x route-reflector-client).

- A route learned from a non-RR client is advertised to RR clients but not to non-RR clients.
- A route learned from a RR client is advertised to both RR clients and non-RR clients. Even the RR client that advertised the route will receive a copy and discards it because it sees itself as the originator.
- A route learned from an EBGP neighbor is advertised to both RR clients and non-RR clients.

Topology

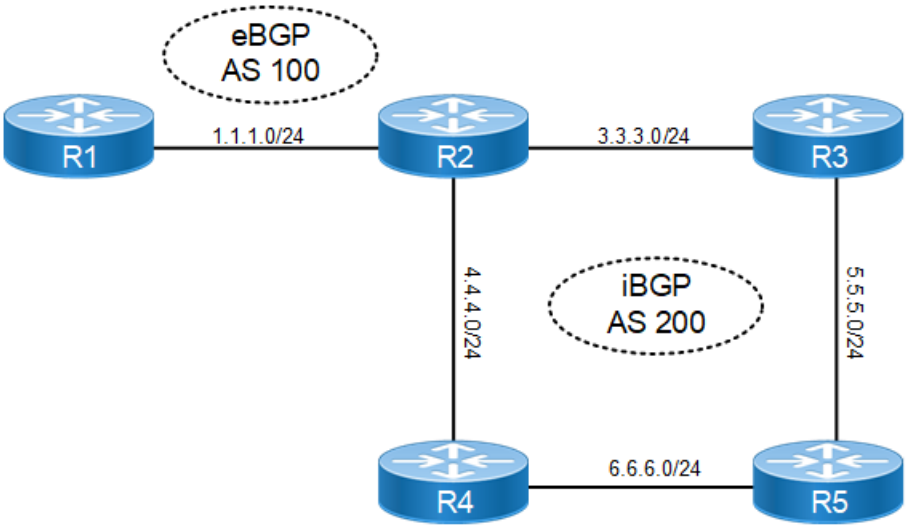

Figure 1-13: eBGP and iBGP Route Reflector Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip addr 1.1.1.1/24`	Specify IP address for the interface.
`(config-if)#exit`	Exit interface mode
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 100.100.100.100/32 secondary`	Specify IP address for the interface.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#router bgp 100`	Define the routing process with AS number `100`.
`(config-router)#neighbor 1.1.1.2 remote-as 200`	Define the eBGP neighbor (R2).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 1.1.1.2 activate`	Activate the neighbor under address family mode
`(config-router-af)#network 100.100.100.100/32`	Advertise a route via eBGP connection to R2.
`(config-router-af)#exit-address-family`	Exit router mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 1.1.1.2/24`	Specify IP address for the interface.
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 3.3.3.2/24`	Specify IP address for the interface.

© 2025 IP Infusion Inc. Proprietary

(config-if)#exit	Exit interface mode
(config)#interface eth3	Enter interface mode
(config-if)#ip address 4.4.4.2/24	Specify IP address for the interface.
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Define the routing process with AS number 200.
(config-router)#neighbor 1.1.1.1 remote-as 100	Define the eBGP neighbor (R1).
(config-router)#neighbor 4.4.4.1 remote-as 200	Define the iBGP neighbor (R4).
(config-router)#neighbor 3.3.3.1 remote-as 200	Define the iBGP neighbor (R3).
(config-router)#bgp cluster-id 4	Define a cluster ID (4) when multiple Route Reflectors exist.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 1.1.1.1 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 4.4.4.1 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 3.3.3.1 activate	Activate the neighbor under address family mode
(config-router-af)#neighbor 3.3.3.1 route-reflector-client	Configure R2 as the Route-Reflector and neighbor R3 as its client.
(config-router-af)#neighbor 4.4.4.1 route-reflector-client	Configure R2 as the Route-Reflector and neighbor R4 as its client.
(config-router-af)#exit-address-family	Exit address-family mode.

R3

#configure terminal	Enter configure mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 3.3.3.1/24	Assign an IP address
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip addr 5.5.5.1/24	Assign an IP address
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Define the routing process with AS number 200.
(config-router)#neighbor 3.3.3.2 remote-as 200	Define the iBGP neighbor (R2).
(config-router)#neighbor 5.5.5.2 remote-as 200	Define the iBGP neighbor (R5).
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 3.3.3.2 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 5.5.5.2 activate	Activate the neighbor under address family mode

(config-router-af)#neighbor 5.5.5.2 route-reflector-client	Configure R3 as the Route-Reflector and neighbor R5 as its client.
(config-router-af)#exit-address-family	Exit address-family mode.

R4

#configure terminal	Enter configure mode
(config)#interface eth2	Enter interface mode
(config-if)#ip address 6.6.6.1/24	Specify an IP address for the interface.
(config-if)#exit	Exit interface mode
(config)#interface eth4	Enter interface mode
(config-if)#ip address 4.4.4.1/24	Specify an IP address for the interface.
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Define the routing process with AS number 200.
(config-router)#neighbor 4.4.4.2 remote-as 200	Define the iBGP neighbor (R2).
(config-router)#neighbor 6.6.6.2 remote-as 200	Define the iBGP neighbor (R5).
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 4.4.4.2 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 6.6.6.2 activate	Activate the neighbor under address family mode
(config-router-af)#exit-address-family	Exit address-family mode.
(config-router)#exit	Exit router mode.

R5

#configure terminal	Enter configure mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 5.5.5.2/24	Specify an IP address for the interface.
(config-if)#exit	Exit interface mode
(config-if)#interface eth2	Enter interface mode
(config-if)#ip address 6.6.6.2/24	Specify an IP address for the interface.
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Define the routing process with AS number 200.
(config-router)#neighbor 5.5.5.1 remote-as 200	Define the iBGP neighbor (R3).
(config-router)#neighbor 6.6.6.1 remote-as 200	Define the iBGP neighbor (R4).
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 5.5.5.1 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 6.6.6.1 activate	Activate the neighbor under address family mode
(config-router-af)#exit-address-family	Exit address-family mode.

Validation

R2

```
#show ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 10.12.4.196, remote router ID 192.160.50.2
  BGP state = Established, up for 00:14:41
  Last read 00:00:11, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 32 messages, 0 notifications, 0 in queue
  Sent 31 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.2, Local port: 50649
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 1.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 3.3.3.1, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.3, remote router ID 192.160.50.4
  BGP state = Established, up for 00:04:17
  Last read 00:00:17, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 10 messages, 0 notifications, 0 in queue
  Sent 13 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 3, Offset 0, Mask 0x8
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 3.3.3.2, Local port: 179
Foreign host: 3.3.3.1, Foreign port: 32973
Nexthop: 3.3.3.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network

BGP neighbor is 4.4.4.1, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.3, remote router ID 192.160.50.6
  BGP state = Established, up for 00:00:16
  Last read 00:00:16, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 2 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 4.4.4.2, Local port: 179
Foreign host: 4.4.4.1, Foreign port: 60398
Nexthop: 4.4.4.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     1.1.1.0/24 is directly connected, eth1, 00:16:10
C     3.3.3.0/24 is directly connected, eth2, 00:15:59
C     4.4.4.0/24 is directly connected, eth3, 00:15:49
B     100.100.100.100/32 [20/0] via 1.1.1.1, eth1, 00:14:53
C     127.0.0.0/8 is directly connected, lo, 00:32:26
C     192.160.50.0/24 is directly connected, eth0, 00:32:22

Gateway of last resort is not set


#show ip bgp
BGP table version is 2, local router ID is 192.160.50.3
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric    LocPrf       Weight Path
*>   100.100.100.100/32
                      1.1.1.1              0         100           0       100
i

Total number of prefixes 1
Total number of neighbors 3
```

R1

```
#show bgp neighbors
BGP neighbor is 1.1.1.2, remote AS 200, local AS 100, external link
  BGP version 4, local router ID 10.12.4.142, remote router ID 10.12.4.196
  BGP state = Established, up for 00:16:11
  Last read 00:00:11, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 34 messages, 0 notifications, 0 in queue
  Sent 36 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 179
Foreign host: 1.1.1.2, Foreign port: 50649
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


#show ip bgp summary
BGP router identifier 192.160.50.2, local AS number 100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1.1.1.2                  4   200   34         36       1      0      0
00:16:18              0

Total number of neighbors 1

Total number of Established sessions 1
#
```

R3

```
#show ip bgp
BGP table version is 1, local router ID is 192.160.50.4
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
* i  100.100.100.100/32
                      1.1.1.1              0         100            0       100
i

Total number of prefixes 1
#

#show ip bgp neighbors
BGP neighbor is 3.3.3.2, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.4, remote router ID 192.160.50.3
  BGP state = Established, up for 00:06:15
  Last read 00:00:15, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 15 messages, 0 notifications, 0 in queue
  Sent 14 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 3.3.3.1, Local port: 32973
Foreign host: 3.3.3.2, Foreign port: 179
Nexthop: 3.3.3.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 5.5.5.2, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.4, remote router ID 192.160.50.5
  BGP state = Established, up for 00:03:35
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 9 messages, 0 notifications, 0 in queue
  Sent 10 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
```

```
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 5.5.5.1, Local port: 179
Foreign host: 5.5.5.2, Foreign port: 39271
Nexthop: 5.5.5.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#

#show ip bgp summary
BGP router identifier 192.160.50.4, local AS number 200
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
3.3.3.2                  4   200   15         14       1      0      0
00:06:26               1
5.5.5.2                  4   200    9         10       1      0      0
00:03:46               0

Total number of neighbors 2
```

R4

```
#show ip bgp
BGP table version is 1, local router ID is 192.160.50.6
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
* i  100.100.100.100/32
                      1.1.1.1              0         100           0       100
i

Total number of prefixes 1
#

#sh ip bgp neighbors
BGP neighbor is 4.4.4.2, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.6, remote router ID 192.160.50.3
  BGP state = Established, up for 00:03:58
  Last read 00:00:28, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 10 messages, 0 notifications, 0 in queue
  Sent 9 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
```

```
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 4.4.4.1, Local port: 60398
Foreign host: 4.4.4.2, Foreign port: 179
Nexthop: 4.4.4.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 6.6.6.2, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.6, remote router ID 192.160.50.5
  BGP state = Established, up for 00:03:52
  Last read 00:00:22, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 9 messages, 0 notifications, 0 in queue
  Sent 9 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 6.6.6.1, Local port: 48257
Foreign host: 6.6.6.2, Foreign port: 179
Nexthop: 6.6.6.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#


#show ip bgp summary
BGP router identifier 192.160.50.6, local AS number 200
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
4.4.4.2                  4   200   11         10       1      0      0
00:04:09              1
6.6.6.2                  4   200   10         10       1      0      0
00:04:03              0

Total number of neighbors 2
```

```
Total number of Established sessions 2
```

R5

```
#show ip bgp neighbors
BGP neighbor is 5.5.5.1, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.5, remote router ID 192.160.50.4
  BGP state = Established, up for 00:09:04
  Last read 00:00:04, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 20 messages, 0 notifications, 0 in queue
  Sent 20 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 5.5.5.2, Local port: 39271
Foreign host: 5.5.5.1, Foreign port: 179
Nexthop: 5.5.5.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 6.6.6.1, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.160.50.5, remote router ID 192.160.50.6
  BGP state = Established, up for 00:07:36
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 17 messages, 0 notifications, 0 in queue
  Sent 18 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 6.6.6.2, Local port: 179
Foreign host: 6.6.6.1, Foreign port: 48257
Nexthop: 6.6.6.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network
#

#sh ip bgp summary
BGP router identifier 192.160.50.5, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V    AS    MsgRcv    MsgSen TblVer    InQ    OutQ    Up/
Down    State/PfxRcd
5.5.5.1                   4    200    20          20        1       0       0
00:09:20                0
6.6.6.1                   4    200    17          18        1       0       0
00:07:52                0

Total number of neighbors 2

Total number of Established sessions 2
#
```

BGP Confederations

In BGP, nodes running iBGP protocols must be interconnected forming a full mesh. Confederation solves the iBGP full-mesh network complexity and inefficiency by splitting a large autonomous system domain into smaller autonomous system domains, called member autonomous systems. Member autonomous systems can form eBGP connections among themselves, to prevent full-mesh connections among each iBGP-running node.

The `bgp confederation identifier` command tells the router that it is a member of a confederation and the confederation ID. The `bgp confederation peers` command lists the member AS to which the router is connected.

In the following example, R1, R2, and R3 are members of the same confederation with different AS numbers.

Topology

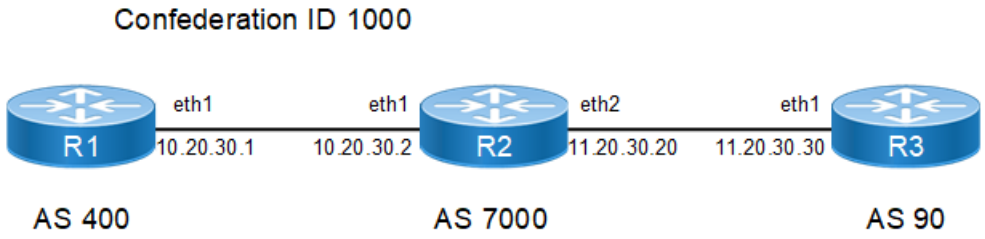

Figure 1-14: BGP Confederation

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 400`	Assign the ASN value (`400`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID, the externally visible autonomous system number that identifies the BGP confederation as a whole.

`#configure terminal`	Enter configure mode.
`(config-router)#bgp confederation peers 7000`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 10.20.30.2 remote-as 7000`	Specify the neighbor's IP address (`10.20.30.2`) and the ASN value of the neighbor (`7000`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.2 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 7000`	Assign the ASN value (`7000`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 400 90`	Specify the neighbor ASN values for confederation membership.
`(config-router)#neighbor 10.20.30.1 remote-as 400`	Specify the neighbor's IP address (`10.20.30.1`) and the ASN value of the neighbor (`400`).
`(config-router)#neighbor 11.20.30.30 remote-as 90`	Specify the neighbor's IP address (`11.20.30.30`) and the ASN value of the neighbor (`90`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.1 activate`	Activate the neighbor under address family mode
`(config-router-af)# neighbor 11.20.30.30 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R3

`#configure terminal`	Enter configure mode.
`(config-router)#router bgp 90`	Assign the ASN value (`90`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 7000`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 11.20.30.20 remote-as 7000`	Specify the neighbor's IP address (`11.20.30.20`) and the ASN value of the neighbor (`7000`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 11.20.30.20 activate`	Exit address-family mode.
`(config-router-af)#exit-address-family`	Exit address-family mode.

Validation

R2

```
#sh ip bgp summary
BGP router identifier 192.168.52.3, local AS number 7000
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.20.30.1               4   400    5          5       1      0      0
00:01:36               0
11.20.30.30              4    90    2          3       1      0      0
00:00:24               0

Total number of neighbors 2

Total number of Established sessions 2

#show ip bgp neighbors
BGP neighbor is 10.20.30.1, remote AS 400, local AS 7000, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  Neighbor under common administration
  BGP state = Established, up for 00:01:25
  Last read 00:01:25, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.20.30.2, Local port: 35108
Foreign host: 10.20.30.1, Foreign port: 179
Nexthop: 10.20.30.2
Nexthop global: fe80::a00:27ff:fe21:7ed2
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.20.30.30, remote AS 90, local AS 7000, external link
  BGP version 4, remote router ID 192.168.56.103
  Neighbor under common administration
  BGP state = Established, up for 00:00:13
  Last read 00:00:13, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
```

```
  Received 2 messages, 0 notifications, 0 in queue
  Sent 3 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.20.30.20, Local port: 179
Foreign host: 11.20.30.30, Foreign port: 33465
Nexthop: 11.20.30.20
Nexthop global: fe80::a00:27ff:fed0:57d1
Nexthop local: ::
BGP connection: non shared network
```

R1

```
#show ip bgp neighbors
BGP neighbor is 10.20.30.2, remote AS 7000, local AS 400, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  Neighbor under common administration
  BGP state = Established, up for 00:01:51
  Last read 00:01:51, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 5 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.20.30.1, Local port: 179
Foreign host: 10.20.30.2, Foreign port: 35108
Nexthop: 10.20.30.1
Nexthop global: fe80::a00:27ff:fe50:6a9b
Nexthop local: ::
BGP connection: non shared network


#sh ip bgp summary
BGP router identifier 192.168.52.3, local AS number 400
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.20.30.2               4  7000    5          6       3      0      0
00:01:57              0

Total number of neighbors 1

Total number of Established sessions 1
```

R3

```
#sh ip bgp neighbors
BGP neighbor is 11.20.30.20, remote AS 7000, local AS 90, external link
  BGP version 4, local router ID 192.168.52.5, remote router ID 192.168.52.3
  Neighbor under common administration
  BGP state = Established, up for 00:00:04
  Last read 00:00:04, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 2 messages, 0 notifications, 0 in queue
  Sent 2 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.20.30.30, Local port: 33465
Foreign host: 11.20.30.20, Foreign port: 179
Nexthop: 11.20.30.30
Nexthop global: fe80::a00:27ff:fe24:5dc9
Nexthop local: ::
BGP connection: non shared network

#sh ip bgp summary
BGP router identifier 192.168.56.103, local AS number 90
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
11.20.30.20              4  7000    3          3       1      0      0
00:00:55              0

Total number of neighbors 1

Total number of Established sessions 1
```

Multiple Autonomous Systems

In the following example, R1 and R2 are members of the same confederation with the same AS numbers, and R3 is a member of the same confederation with a different AS number.

Topology

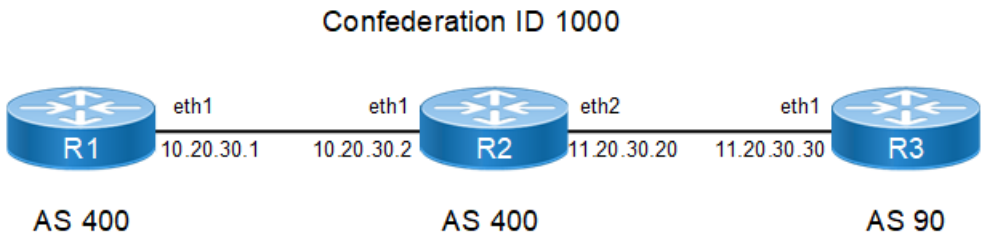

Figure 1-15: BGP Confederation with Multiple AS

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 400`	Assign the ASN value (`400`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#neighbor 10.20.30.2 remote-as 400`	Specify the neighbor's IP address (`10.20.30.2`) and the ASN value of the neighbor (`400`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.2 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 400`	Assign the ASN value (`400`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 90`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 10.20.30.1 remote-as 400`	Specify the neighbor's IP address (`10.20.30.1`) and the ASN value of the neighbor (`400`).
`(config-router)#neighbor 11.20.30.30 remote-as 90`	Specify the neighbor's IP address (`11.20.30.30`) and the ASN value of the neighbor (`90`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.1 activate`	Activate the neighbor under address family mode
`(config-router-af)# neighbor 11.20.30.30 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R3

`#configure terminal`	Enter configure mode.
`(config)#router bgp 90`	Assign the ASN value (`90`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 400`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 11.20.30.20 remote-as 400`	Specify the neighbor's IP address (`11.20.30.20`) and the ASN value of the neighbor (`400`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 11.20.30.20 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

Validation

R2

```
#show ip bgp summary
BGP router identifier 192.168.52.3, local AS number 400
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.20.30.1               4   400   16         16       1      0      0
00:07:27              0
11.20.30.30              4    90   32         42       1      0      0
00:00:27              0

Total number of neighbors 2

Total number of Established sessions 2
#show ip bgp neighbors
BGP neighbor is 10.20.30.1, remote AS 400, local AS 400, internal link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  BGP state = Established, up for 00:08:10
  Last read 00:08:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 18 messages, 0 notifications, 0 in queue
  Sent 18 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
```

```
 Connections established 1; dropped 0
Local host: 10.20.30.2, Local port: 35214
Foreign host: 10.20.30.1, Foreign port: 179
Nexthop: 10.20.30.2
Nexthop global: fe80::a00:27ff:fe21:7ed2
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.20.30.30, remote AS 90, local AS 400, external link
  BGP version 4, remote router ID 192.168.56.103
  Neighbor under common administration
  BGP state = Established, up for 00:01:10
  Last read 00:01:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 20 messages, 14 notifications, 0 in queue
  Sent 42 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.20.30.20, Local port: 179
Foreign host: 11.20.30.30, Foreign port: 33623
Nexthop: 11.20.30.20
Nexthop global: fe80::a00:27ff:fed0:57d1
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:01:36, due to BGP Notification sent
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

R1

```
#show ip bgp neighbors
BGP neighbor is 10.20.30.2, remote AS 400, local AS 400, internal link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  BGP state = Established, up for 00:08:41
  Last read 00:08:41, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 34 messages, 0 notifications, 0 in queue
  Sent 35 messages, 3 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 16, neighbor version 16
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
```

```
  0 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
Local host: 10.20.30.1, Local port: 179
Foreign host: 10.20.30.2, Foreign port: 35214
Nexthop: 10.20.30.1
Nexthop global: fe80::a00:27ff:fe50:6a9b
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:09:03, due to BGP Notification sent
Notification Error Message: (OPEN Message Error/Bad Peer AS.)

#show ip bgp summary
BGP router identifier 192.168.52.2, local AS number 400
BGP table version is 16
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.20.30.2               4   400   34         38      16      0      0
00:08:44              0

Total number of neighbors 1

Total number of Established sessions 1
```

R3

```
#show ip bgp summary
BGP router identifier 192.168.52.5, local AS number 90
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
11.20.30.20              4   400    2          2       1      0      0
00:00:15              0

Total number of neighbors 1

Total number of Established sessions 1

#show ip bgp neighbors
BGP neighbor is 11.20.30.20, remote AS 400, local AS 90, external link
  BGP version 4, local router ID 192.168.52.5, remote router ID 192.168.52.3
  Neighbor under common administration
  BGP state = Established, up for 00:02:24
  Last read 00:02:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 6 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
```

```
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.20.30.30, Local port: 33623
Foreign host: 11.20.30.20, Foreign port: 179
Nexthop: 11.20.30.30
Nexthop global: fe80::a00:27ff:fe24:5dc9
Nexthop local: ::
BGP connection: non shared network
```

Outside Autonomous System

In the following example, R1 and R2 are members of the same confederation with different AS numbers, and R3 is a member outside the confederation.

Topology

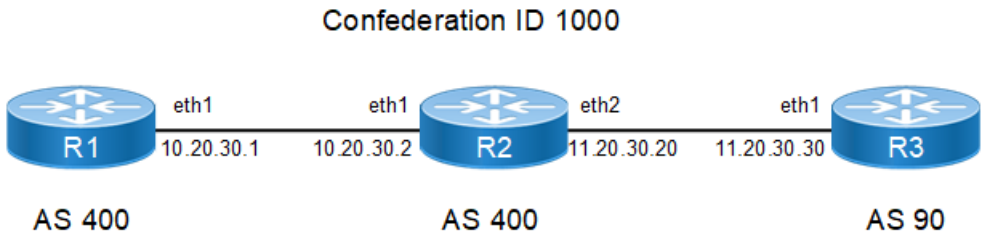

Figure 1-16: Single Confederation with Outside AS

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 400`	Assign the ASN value (`400`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 7000`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 10.20.30.2 remote-as 7000`	Specify the neighbor's IP address (`10.20.30.2`) and the ASN value of the neighbor (`7000`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.2 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 7000`	Assign the ASN value (`7000`) to the router.
`(config-router)#bgp confederation identifier 1000`	Specify the BGP confederation ID.
`(config-router)#bgp confederation peers 400`	Specify the neighbor ASN value for confederation membership.
`(config-router)#neighbor 10.20.30.1 remote-as 400`	Specify the neighbor's IP address (`10.20.30.1`) and the ASN value of the neighbor (`400`).
`(config-router)#neighbor 11.20.30.30 remote-as 90`	Specify the neighbor's IP address (`11.20.30.30`) and the ASN value of the neighbor (`90`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.1 activate`	Activate the neighbor under address family mode
`(config-router-af)# neighbor 11.20.30.30 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

R3

`#configure terminal`	Enter configure mode.
`(config)#router bgp 90`	Assign the ASN value (`90`) to the router.
`(config-router)#neighbor 11.20.30.20 remote-as 1000`	Specify the neighbor's IP address (`11.20.30.20`) and the BGP confederation ID (`1000`).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 11.20.30.20 activate`	Activate the neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address-family mode.

Validation

R3

```
#show ip bgp neighbors
BGP neighbor is 11.20.30.20, remote AS 1000, local AS 90, external link
  BGP version 4, local router ID 192.168.52.5, remote router ID 192.168.52.3
  BGP state = Established, up for 00:01:10
  Last read 00:01:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 112 messages, 1 notifications, 0 in queue
  Sent 142 messages, 88 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
```

```
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
Local host: 11.20.30.30, Local port: 33951
Foreign host: 11.20.30.20, Foreign port: 179
Nexthop: 11.20.30.30
Nexthop global: fe80::a00:27ff:fe24:5dc9
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:01:26, due to BGP Notification sent
Notification Error Message: (OPEN Message Error/Bad Peer AS.)

#sh ip bgp summary
BGP router identifier 192.168.52.5, local AS number 90
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
11.20.30.20              4  1000  113        230       1      0      0
00:01:13              0

Total number of neighbors 1

Total number of Established sessions 1
```

R2

```
#show ip bgp summary
BGP router identifier 192.168.52.3, local AS number 7000
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
10.20.30.1               4   400   22         22       1      0      0
00:10:04              0
11.20.30.30              4    90  179        202       1      0      0
00:00:42              0

Total number of neighbors 2

Total number of Established sessions 2

#show ip bgp neighbors
BGP neighbor is 10.20.30.1, remote AS 400, local AS 7000, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.3
  Neighbor under common administration
  BGP state = Established, up for 00:11:06
  Last read 00:11:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
```

```
  Received 24 messages, 0 notifications, 0 in queue
  Sent 24 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.20.30.2, Local port: 35444
Foreign host: 10.20.30.1, Foreign port: 179
Nexthop: 10.20.30.2
Nexthop global: fe80::a00:27ff:fe21:7ed2
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.20.30.30, remote AS 90, local AS 1000, external link
  BGP version 4, remote router ID 192.168.56.103
  BGP state = Established, up for 00:01:44
  Last read 00:01:44, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 93 messages, 88 notifications, 0 in queue
  Sent 204 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.20.30.20, Local port: 179
Foreign host: 11.20.30.30, Foreign port: 33951
Nexthop: 11.20.30.20
Nexthop global: fe80::a00:27ff:fed0:57d1
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:02:00, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

R1

```
#sh ip bgp summary
BGP router identifier 192.168.52.2, local AS number 400
BGP table version is 34
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
```

```
10.20.30.2                4  7000   77          91       34       0       0
00:10:18               0

Total number of neighbors 1

Total number of Established sessions 1
#sh ip bgp neighbors
BGP neighbor is 10.20.30.2, remote AS 7000, local AS 400, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  Neighbor under common administration
  BGP state = Established, up for 00:11:40
  Last read 00:11:40, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 80 messages, 0 notifications, 0 in queue
  Sent 82 messages, 12 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 35, neighbor version 35
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 3; dropped 2
Local host: 10.20.30.1, Local port: 179
Foreign host: 10.20.30.2, Foreign port: 35444
Nexthop: 10.20.30.1
Nexthop global: fe80::a00:27ff:fe50:6a9b
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:12:47, due to BGP Notification sent
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

Dynamic BGP Peering

BGP Dynamic Neighbors is a quick way of setting up BGP on device like a Hub router where user is expecting numerous BGP neighbors. Before dynamic neighbors, user had to provide a large amount of configuration to work with all these neighbors. This new feature dramatically reduces the amount and complexity of CLI configuration on the router and save CPU and memory usage.

BGP dynamic neighbor support allows BGP peering to a group of remote neighbors that are defined by a range of IP addresses. Each range can be configured as a subnet IP address. BGP dynamic neighbors are configured using a range of IP addresses and BGP peer groups.

After a subnet range is configured for a BGP peer group and a TCP session is initiated by another router for an IP address in the subnet range, a new BGP neighbor is dynamically created as a member of that group. After the initial configuration of subnet ranges and activation of the peer group, dynamic BGP neighbor creation does not require any further CLI configuration on the initial router. Other routers can establish a BGP session with the initial router, but the initial router need not establish a BGP session to other routers if the IP address of the remote peer used for the BGP session is not within the configured range.

A dynamic BGP neighbor will inherit any configuration for the peer group. In larger BGP networks, implementing BGP dynamic neighbors can reduce the amount and complexity of CLI configuration and save CPU and memory usage. Both IPv4 and IPV6 peering is supported.

IPv4 IBGP Configuration

Below figure displays the minimum configuration required to enable BGP on an interface. R1 and R2 are two routers belonging to the same AS, AS100, connecting to network 11.11.11.0/24 and 11.11.12.0/24. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

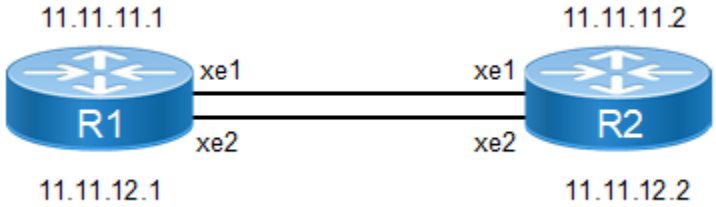

Figure 1-17: IPv4 IBGP Peering

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ip address 33.33.33.33/32 secondary`	Assign a secondary IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)#ip address 11.11.11.1/24`	Assign IP address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode for xe2.
`(config-if)#ip add 11.11.12.1/24`	Assign IP address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode.
`(config-router)#bgp router-id 1.1.1.1`	Assign a BGP router ID.
`(config-router)#neighbor IPV4_IBGP_PEER peer- group range 11.11.0.0/16`	Create a dynamic peer-group, IPV4_IBGP_PEER with a dynamic range 11.11.0.0/16
`(config-router)#neighbor IPV4_IBGP_PEER remote-as 100`	Assign a remote AS for the peer-group, IPV4_IBGP_PEER.
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor IPV4_IBGP_PEER activate`	Activate peer-group in the address family
`(config-router-af)#network 33.33.33.33/32`	Advertise the loopback network into BGP.
`(config-router-af)#end`	Exit from Router BGP mode.

R2

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ip address 22.22.22.22/32 secondary	Assign a secondary IP address.
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 11.11.11.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ip add 11.11.12.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode.
(config-router)#bgp router-id 2.2.2.2	Assign a BGP router ID.
(config-router)#neighbor 11.11.11.1 remote-as 100	Create a static BGP neighbor 11.11.11.1 in remote AS 100.
(config-router)#neighbor 11.11.12.1 remote-as 100	Create a static BGP neighbor 11.11.12.1 in remote AS 100.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 11.11.12.1 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 11.11.12.1 activate	Activate the neighbor under address family mode
(config-router-af)# network 22.22.22.22/32	Advertise the loopback network into BGP.
(config-router-af)#exit-address-family	Exit address-family mode.

Validation

R1

```
#show ip bgp summary

BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11.11.11.2              4   100   42         43       2      0      0
00:20:25              1
*11.11.12.2              4   100   42         43       2      0      0
00:20:25              1
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV4_IBGP_PEER
  listen range:  11.11.0.0/16
```

```
  Total number of dynamically created neighbors/limit: 2/(200)

Total number of dynamically created neighbors: 2
Total number of activated dynamic peer-groups for IPv4 Unicast address-family:
1

Total number of neighbors 2
Total number of Established sessions 2


#show ip bgp neighbors

BGP neighbor is 11.11.11.2, remote AS 100, local AS 100, internal link
 Member of peer-group IPV4_IBGP_PEER for session parameters
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 00:21:56
  Last read 00:00:27, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 46 messages, 0 notifications, 0 in queue
  Sent 46 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 0, Offset 0, Mask 0x1
  IPV4_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.11.1, Local port: 40361
Foreign host: 11.11.11.2, Foreign port: 179
Nexthop: 11.11.11.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.11.12.2, remote AS 100, local AS 100, internal link
 Member of peer-group IPV4_IBGP_PEER for session parameters
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 00:21:56
  Last read 00:00:27, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 46 messages, 0 notifications, 0 in queue
  Sent 46 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  IPV4_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
```

```
  1 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.12.1, Local port: 33478
Foreign host: 11.11.12.2, Foreign port: 179
Nexthop: 11.11.12.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show running-config bgp
router bgp 100
bgp router-id 1.1.1.1 network 33.33.33.33/32
neighbor IPV4_IBGP_PEER peer-group range 11.11.0.0/16 neighbor IPV4_IBGP_PEER
remote-as 100
!
address-family ipv4 unicast
 neighbor IPV4_IBGP_PEER activate
 exit-address-family

#show ip bgp
BGP table version is 2, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

NetworkNext HopMetricLocPrfWeight Path
*>i22.22.22.22/3211.11.11.201000i
* i11.11.12.201000i
*>33.33.33.33/320.0.0.0010032768i

Total number of prefixes 2
```

R2

```
#show ip bgp summary
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
11.11.11.1               4   100   55         56       3      0      0
00:26:21              1
11.11.12.1               4   100   55         56       3      0      0
00:26:21              1

Total number of neighbors 2
Total number of Established sessions 2


#show bgp neighbors
BGP neighbor is 11.11.11.1, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
```

```
  BGP state = Established, up for 00:26:43
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 56 messages, 0 notifications, 0 in queue
  Sent 57 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.11.2, Local port: 179
Foreign host: 11.11.11.1, Foreign port: 40361
Nexthop: 11.11.11.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.11.12.1, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:26:43
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 56 messages, 0 notifications, 0 in queue
  Sent 57 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.12.2, Local port: 179
Foreign host: 11.11.12.1, Foreign port: 33478
Nexthop: 11.11.12.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip bgp
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
    Network          Next Hop            Metric    LocPrf       Weight Path
*>   22.22.22.22/32   0.0.0.0              0         100           32768    i
*>i  33.33.33.33/32   11.11.11.1            0         100           0        i
* i                   11.11.12.1            0         100           0        i

Total number of prefixes 2
```

IPv4 IBGP VRF Configuration

Below figure displays the minimum configuration required to enable BGP on an interface with vrf enabled on the device and interface being part of vrf. R1 and R2 are two routers belonging to the same AS, AS100, connecting to network 11.11.11.0/24 and 11.11.12.0/24. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

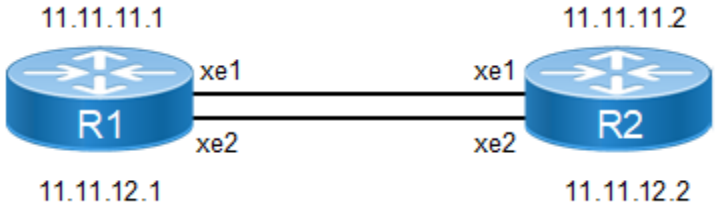

Figure 1-18: IPv4 VRF IBGP Peering

R1

`#Configure terminal`	Enter Configuration mode.
`(config)#ip vrf vrfA`	Create a VRF, vrfA on the router.
`(config-vrf)#rd 1:1`	Assign a route distinguisher to VRF.
`(config-if)#exit`	Exit VRF mode and return to Configure mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)# ip vrf forwarding vrfA`	Assign IP address to VRF.
`(config-if)#ip address 11.11.11.1/24`	Assign IP address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode for xe2.
`(config-if)# ip vrf forwarding vrfA`	Assign IP address to VRF.
`(config-if)#ip add 11.11.12.1/24`	Assign IP address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode.
`(config-router)#bgp router-id 1.1.1.1`	Assign a BGP router ID.
`(config-router)#address-family ipv4 vrf vrfA`	Enter IPv4 VRF Address Family.
`(config-router-af)#neighbor IPV4_IBGP_PEER peer-group range 11.11.0.0/16`	Create a dynamic peer-group, IPV4_IBGP_PEER with a dynamic range 11.11.0.0/16.
`(config-router-af)#neighbor IPV4_IBGP_PEER remote-as 100`	Assign a remote AS for the peer-group, IPV4_IBGP_PEER.

(config-router-af)# neighbor IPV4_IBGP_PEER activate	Activate peer-group in the address family
(config-router-af)#network 33.33.33.33/32	Advertise the loopback network into BGP
(config-router-af)#end	Exit from Router BGP mode.

R2

#configure terminal	Enter configure mode.
(config)#ip vrf vrfA	Create a VRF, vrfA on router.
(config-vrf)#rd 2:1	Assign a route distinguisher to VRF.
(config-if)#exit	Exit VRF mode and return to Configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)# ip vrf forwarding vrfA	Assign IP address to VRF.
(config-if)#ip address 11.11.11.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)# ip vrf forwarding vrfA	Assign IP address to VRF.
(config-if)#ip add 11.11.12.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode.
(config-router)#bgp router-id 2.2.2.2	Assign a BGP router ID.
(config-router)#address-family ipv4 vrf vrfA	Enter IPv4 VRF Address Family.
(config-router-af)#neighbor 11.11.11.1 remote-as 100	Create a static BGP neighbor 11.11.11.1 in remote AS 100.
(config-router-af)# neighbor 11.11.11.1 activate	Activate neighbor in the address family
(config-router-af)#neighbor 11.11.12.1 remote-as 100	Create a static BGP neighbor 11.11.12.1 in remote AS 100
(config-router-af)# neighbor 11.11.12.1 activate	Activate neighbor in the address family
(config-router-af)#network 22.22.22.22/32	Advertise the loopback network into BGP.
(config-router-af)#end	Exit from Router BGP mode.

Validation

R1

```
#show running-config bgp
!
router bgp 100
 bgp router-id 1.1.1.1
 !
 address-family ipv4 vrf vrfA
 neighbor IPV4_VRF_IBGP_PEER peer-group range 11.11.0.0/16
 neighbor IPV4_VRF_IBGP_PEER remote-as 100
 neighbor IPV4_VRF_IBGP_PEER activate
```

```
 network 33.33.33.33/32
 exit-address-family
!

#show ip bgp summary vrf vrfA
BGP router identifier 11.11.11.1, local AS number 100
BGP VRF vrfA Route Distinguisher: 1:1
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11.11.11.2              4   100    3          3       1      0      0
00:01:00              0
*11.11.12.2              4   100    3          3       1      0      0
00:00:55              0
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV4_IBGP_PEER
  listen range:  11.11.0.0/16
  Total number of dynamically created neighbors/limit: 2/(200)

Total number of dynamically created neighbors: 2
Total number of activated dynamic peer-groups for IPv4 Unicast address-family:
1

Total number of neighbors 2

Total number of Established sessions 2


#show bgp neighbors
BGP neighbor is 11.11.11.2, vrf vrfA, remote AS 100, local AS 100, internal
link
 Member of peer-group IPV4_IBGP_PEER for session parameters
  BGP version 4, local router ID 11.11.11.1, remote router ID 11.11.11.2
  BGP state = Established, up for 00:07:26
  Last read 00:00:26, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 16 messages, 0 notifications, 0 in queue
  Sent 16 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  IPV4_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.11.1, Local port: 36365
Foreign host: 11.11.11.2, Foreign port: 179
Nexthop: 11.11.11.1
```

```
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.11.12.2, vrf vrfA, remote AS 100, local AS 100, internal
link
 Member of peer-group IPV4_IBGP_PEER for session parameters
  BGP version 4, local router ID 11.11.11.1, remote router ID 11.11.11.2
  BGP state = Established, up for 00:07:21
  Last read 00:00:21, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 16 messages, 0 notifications, 0 in queue
  Sent 16 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  IPV4_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.12.1, Local port: 38144
Foreign host: 11.11.12.2, Foreign port: 179
Nexthop: 11.11.12.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

IPv4 EBGP Configuration

Below figure displays the minimum configuration required to enable BGP on an interface. R1, R2 and R3 are three routers belonging to the different AS, AS100 AS200 and AS300, connecting to network 11.11.11.0/24 and 11.11.12.0/24. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

Figure 1-19: IPv4 EBGP Peering

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ip address 33.33.33.33/32 secondary`	Assign a secondary IP address.

(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 11.11.11.1/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode.
(config-router)#bgp router-id 1.1.1.1	Assign a BGP router ID.
(config-router)#neighbor 11.11.11.2 remote-as 200	Create a static neighbor 11.11.11.2 with remote AS 200.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode.
(config-router-af)# neighbor 11.11.11.2 activate	Activate the neighbor under address family mode
(config-router-af)#exit-address-family	Exit address-family mode.

R2

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ip address 22.22.22.22/32 secondary	Assign a secondary IP address.
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 11.11.11.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ip add 11.11.12.2/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 200	Enter Router BGP mode.
(config-router)#bgp router-id 2.2.2.2	Assign a BGP router ID.
(config-router)#neighbor IPV4_EBGP_PEER peer-group range 11.11.0.0/16	Create a dynamic peer-group, IPV4_EBGP_PEER.
(config-router)#neighbor IPV4_EBGP_PEER remote-as 100	Assign remote AS with the peer-group IPV4_EBGP_PEER.
(config-router)#neighbor IPV4_EBGP_PEER optional-as 300	Assign optional AS with the peer-group IPV4_EBGP_PEER
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor IPV4_EBGP_PEER activate	Activate the peer-group in address family
(config-router-af)#network 22.22.22.22/32	Advertise the loopback network into BGP.
(config-router-af)#end	Exit from address family mode.

R3

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode for loopback.

(config-if)#ip address 44.44.44.44/32 secondary	Assign a secondary IP address.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ip add 11.11.12.3/24	Assign IP address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 300	Enter Router BGP mode.
(config-router)#bgp router-id 3.3.3.3	Assign a BGP router ID.
(config-router)#neighbor 11.11.12.2 remote-as 200	Create a static BGP neighbor 11.11.12.2 with remote AS 200.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 11.11.12.2 activate	Activate the neighbor under address family mode
(config-router-af)#network 44.44.44.44/32	Advertise the loopback network into BGP.
(config-router-af)#exit-address-family	Exit address-family mode.

Validation

R2

```
#show ip bgp summary
BGP router identifier 2.2.2.2, local AS number 200
BGP table version is 3
2 BGP AS-PATH entries
0 BGP community entries
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11.11.11.1              4   100   29         29       3      0      0
00:13:10              1
*11.11.12.3              4   300   27         27       3      0      0
00:12:20              1
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV4_EBGP_PEER
  listen range:  11.11.0.0/16
  Total number of dynamically created neighbors/limit: 2/(200)
Total number of dynamically created neighbors: 2
Total number of activated dynamic peer-groups for IPv4 Unicast address-family:
1
Total number of neighbors 2
Total number of Established sessions 2

#show running-config bgp
!
router bgp 200
 bgp router-id 2.2.2.2
 network 22.22.22.22/32
 neighbor IPV4_EBGP_PEER peer-group range 11.11.0.0/16
 neighbor IPV4_EBGP_PEER remote-as 100
 neighbor IPV4_EBGP_PEER optional-as 300
!
 address-family ipv4 unicast
```

```
 neighbor IPV4_EBGP_PEER activate
 network 33.33.33.33/32
 exit-address-family
!


#show bgp neighbors
BGP neighbor is 11.11.11.1, remote AS 100, local AS 200, external link
 Member of peer-group IPV4_EBGP_PEER for session parameters
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:17:15
  Last read 00:00:15, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 37 messages, 0 notifications, 0 in queue
  Sent 38 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  IPV4_EBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.11.2, Local port: 42252
Foreign host: 11.11.11.1, Foreign port: 179
Nexthop: 11.11.11.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.11.12.3, remote AS 300, local AS 200, external link
 Member of peer-group IPV4_EBGP_PEER for session parameters
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 00:13:17
  Last read 00:00:17, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 29 messages, 0 notifications, 0 in queue
  Sent 30 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  IPV4_EBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 11.11.12.2, Local port: 59839
```

```
Foreign host: 11.11.12.3, Foreign port: 179
Nexthop: 11.11.12.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

IPv6 IBGP Configuration

Below figure displays the minimum configuration required to enable BGP on an interface. R1 and R2 are two routers belonging to the same AS, AS100, connecting to network 11:11:11::1/64 and 11:11:12::1/64. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

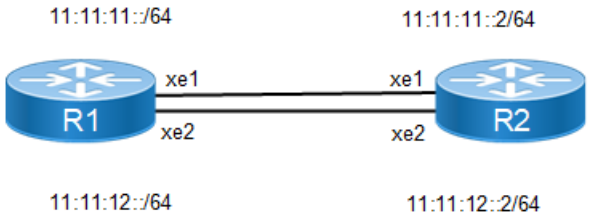

Figure 1-20: IPv6 IBGP Peering

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ipv6 address 33::1/128`	Assign an IPv6 address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)#ipv6 address 11:11:11::1/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode for xe2.
`(config-if)#ipv6 address 11:11:12::1/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode.
`(config-router)#bgp router-id 1.1.1.1`	Assign a BGP router ID.
`(config-router)#neighbor IPV6_IBGP_PEER peer-group range 11:11::/16`	Create a dynamic peer-group, IPV6_IBGP_PEER with a dynamic range 11:11::/16
`(config-router)#neighbor IPV6_IBGP_PEER remote-as 100`	Configure a remote AS with the peer group, IPV6_IBGP_PEER.
`(config-router)#neighbor IPV6_IBGP_PEER limit 1`	Set peer group neighbors limit to 1. Only one BGP session will be up.
`(config-router)#address-family ipv6 unicast`	Enter the IPv6 Unicast Address Family.
`(config-router-af)#neighbor IPV6_IBGP_PEER activate`	Activate the peer group, IPV6_IBGP_PEER in the IPv6 address family.

#configure terminal	Enter configure mode.
(config-router-af)#network 33::1/128	Advertise the loopback network into the BGP IPv6 address family.
(config-router-af)#end	Exit from address family mode.

R2

#Configure terminal	Enter Configuration mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ipv6 address 22::2/128	Assign an IPv6 address.
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ipv6 address 11:11:11::2/64	Assign an IPv6 address to the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ipv6 address 11:11:12::2/64	Assign an IPv6 address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode.
(config-router)#bgp router-id 2.2.2.2	Assign a BGP router ID.
(config-router)#neighbor 11:11:11::1 remote-as 100	Configure BGP neighbor by specifying the neighbor IP address.
(config-router)#neighbor 11:11:12::1 remote-as 100	Configure BGP neighbor by specifying the neighbor IP address.
(config-router)#address-family ipv4 unicast	Enter the Ipv4 Unicast Address Family.
(config-router-af)# neighbor 11:11:11::1 activate	Activate the neighbor under address family mode
(config-router-af)# neighbor 11:11:12::1 activate	Activate the neighbor under address family mode
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#address-family ipv6 unicast	Enter the IPv6 Unicast Address Family.
(config-router-af)#network 22::2/128	Advertise the loopback network into BGP.
(config-router-af)#neighbor 11:11:12::1 activate	Activate the neighbor in the IPv6 address family.
(config-router-af)#neighbor 11:11:11::1 activate	Activate the neighbor in the IPv6 address family.
(config-router-af)#end	Exit from address family mode.

Validation

R1

```
#show ipv6 bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 2
1 BGP AS-PATH entries
```

```
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11:11:11::2             4   100    6          6       2      0      0
00:01:41               1
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV6_IBGP_PEER
  listen range:  11::/16
  Total number of dynamically created neighbors/limit: 1/(1)

Total number of dynamically created neighbors: 1
Total number of activated dynamic peer-groups for IPv6 Unicast address-family:
1

Total number of neighbors 1

Total number of Established sessions 1

#show ip bgp peer-group IPV6_IBGP_PEER

BGP dynamic peer-group is IPV6_IBGP_PEER, IBGP, remote AS 100
  BGP dynamic peer-group IPV6_IBGP_PEER listen range group members:
  11::/16
  BGP version 4
Minimum time between advertisement runs is 5 seconds
For address family: IPv4 Unicast
  Peer-group member:
  *11:11:11::2
  Index 1, Offset 0, Mask 0x2
  0 accepted prefixes, 0 announced prefixes
For address family: IPv6 Unicast
  Peer-group member:
  *11:11:11::2
  Index 0, Offset 0, Mask 0x0
  1 accepted prefixes, 1 announced prefixes


#show bgp ipv6
BGP table version is 2, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  22::2/128        11:11:11::2(fe80::5054:ff:fe95:85ec)
                                          0         100           0        i
*>   33::1/128        ::                   0         100           32768    i

Total number of prefixes 2

#show running-config bgp
!
router bgp 100
 bgp router-id 1.1.1.1
```

```
 neighbor IPV6_IBGP_PEER peer-group range 11::/16
 neighbor IPV6_IBGP_PEER remote-as 100
 neighbor IPV6_IBGP_PEER limit 1
 !
 address-family ipv6 unicast
 network 33::1/128
 neighbor IPV6_IBGP_PEER activate
 exit-address-family


#show bgp neighbors
BGP neighbor is 11:11:11::2, remote AS 100, local AS 100, internal link
 Member of peer-group IPV6_IBGP_PEER for session parameters
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 00:04:17
  Last read 00:00:18, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 0, Offset 0, Mask 0x1
  IPV6_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 2, neighbor version 2
  Index 0, Offset 0, Mask 0x0
  IPV6_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 11:11:11::1, Local port: 42410
Foreign host: 11:11:11::2, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: 11:11:11::1
Nexthop local: fe80::5054:ff:fe51:f74
BGP connection: shared network
```

IPV6 IBGP VRF Configuration

Below figure displays the minimum configuration required to enable BGP on an interface with VRF enabled on the device and interface being part of VRF. R1 and R2 are two routers belonging to the same AS, AS100, connecting to network 11:11:11::1 and 11:11:12::1. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

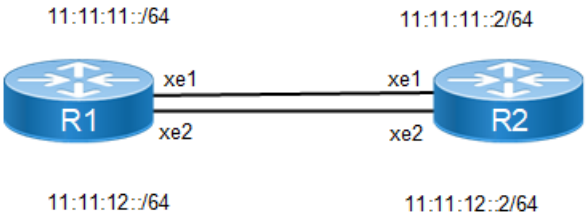

Figure 1-21: IPv6 VRF IBGP peering

R1

`#Configure terminal`	Enter Configuration mode.
`(config)#ip vrf vrfA`	Configure a VRF, vrfA.
`(config-vrf)#rd 1:1`	Configure a route distinguisher to VRF.
`(config-vrf)#router-id 7.7.7.7`	Configure a router ID.
`(config-vrf)#exit`	Exit from VRF mode and return to Configuration mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)# ip vrf forwarding vrfA`	Assign IP address to VRF.
`(config-if)#ipv6 address 11:11:11::1/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode for xe2.
`(config-if)# ip vrf forwarding vrfA`	Assign IP address to VRF.
`(config-if)#ipv6 address 11:11:12::1/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode.
`(config-router)#address-family ipv6 vrf vrfA`	Enter IPv6 VRF Address Family.
`(config-router-af)#neighbor IPV6_VRF_IBGP_PEER peer-group range 11:11::1/16`	Configure a dynamic peer group, IPV6_IBGP_PEER with a dynamic range value.
`(config-router-af)#neighbor IPV6_VRF_IBGP_PEER remote-as 100`	Configure a remote AS with the peer group, IPV6_IBGP_PEER.
`(config-router-af)# neighbor IPV6_VRF_IBGP_PEER activate`	Activate neighbor in the address family
`(config-router-af)#end`	Exit from address family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrfA`	Configure a VRF, vrfA.
`(config-vrf)#rd 2:1`	Configure a route distinguisher to VRF.
`(config-vrf)#router-id 1.1.1.1`	Configure a router ID.
`(config-vrf)#exit`	Exit from VRF mode and return to Configuration mode.
`(config)#interface xe1`	Enter interface mode for xe1.

© 2025 IP Infusion Inc. Proprietary

#configure terminal	Enter configure mode.
(config-if)# ip vrf forwarding vrfA	Assign IP address to VRF.
(config-if)#ipv6 address 11:11:11::2/64	Assign an IPv6 address to the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)# ip vrf forwarding vrfA	Assign IP address to vrf.
(config-if)#ipv6 address 11:11:12::2/64	Assign an IPv6 address to the interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode
(config-router)#address-family ipv6 vrf vrfA	Enter IPv6 VRF Address Family.
(config-router-af)#neighbor 11:11:12::1 remote-as 100	Configure BGP neighbor by specifying the neighbor IP address.
(config-router-af)#neighbor 11:11:11::1 remote-as 100	Configure BGP neighbor by specifying the neighbor IP address.
(config-router-af)#end	Exit from address family mode..

Validation

R1

```
#show ipv6 bgp summary vrf vrfA
BGP router identifier 7.7.7.7, local AS number 100
BGP VRF vrfA Route Distinguisher: 1:1
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11:11:11::2             4   100    6          6       1      0      0
00:00:17              0
*11:11:12::2             4   100    7         10       1      0      0
00:00:15              0
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV6_VRF_IBGP_PEER
  listen range:  11::/16
  Total number of dynamically created neighbors/limit: 2/(200)

Total number of dynamically created neighbors: 2
Total number of activated dynamic peer-groups for IPv6 Unicast address-family:
1

Total number of neighbors 2

Total number of Established sessions 2


#show running-config bgp
!
router bgp 100
```

```
 !
 address-family ipv6 vrf vrfA
 neighbor IPV6_VRF_IBGP_PEER peer-group range 11::/16
 neighbor IPV6_VRF_IBGP_PEER remote-as 100
 neighbor IPV6_VRF_IBGP_PEER activate
 exit-address-family
!

#show ip bgp peer-group vrf vrfA

BGP dynamic peer-group is IPV6_VRF_IBGP_PEER, IBGP, remote AS 100
  BGP dynamic peer-group IPV6_VRF_IBGP_PEER listen range group members:
  11::/16
  BGP version 4
Minimum time between advertisement runs is 5 seconds
For address family: IPv6 Unicast
  Peer-group member:
  *11:11:12::2
  Index 1, Offset 0, Mask 0x2
  0 accepted prefixes, 0 announced prefixes
  Peer-group member:
  *11:11:11::2
  Index 2, Offset 0, Mask 0x4
  0 accepted prefixes, 0 announced prefixes

#show running-config bgp
!
router bgp 100
 !
 address-family ipv6 vrf vrfA
 neighbor IPV6_VRF_IBGP_PEER peer-group range 11::/16
 neighbor IPV6_VRF_IBGP_PEER remote-as 100
 neighbor IPV6_VRF_IBGP_PEER activate
 exit-address-family
!
#show bgp ipv6 neighbors
BGP neighbor is 11:11:11::2, vrf vrfA, remote AS 100, local AS 100, internal
link
 Member of peer-group IPV6_VRF_IBGP_PEER for session parameters
  BGP version 4, local router ID 7.7.7.7, remote router ID 1.1.1.1
  BGP state = Established, up for 00:02:13
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv6 Unicast: advertised and received
  Received 8 messages, 2 notifications, 0 in queue
  Sent 10 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  IPV6_VRF_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
```

```
 Connections established 1; dropped 0
Local host: 11:11:11::1, Local port: 179
Foreign host: 11:11:11::2, Foreign port: 48206
Nexthop: 7.7.7.7
Nexthop global: 11:11:11::1
Nexthop local: fe80::5054:ff:fe51:f74
BGP connection: shared network
Last Reset: 00:02:18, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier.)

BGP neighbor is 11:11:12::2, vrf vrfA, remote AS 100, local AS 100, internal
link
 Member of peer-group IPV6_VRF_IBGP_PEER for session parameters
  BGP version 4, local router ID 7.7.7.7, remote router ID 1.1.1.1
  BGP state = Established, up for 00:02:11
  Last read 00:00:12, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv6 Unicast: advertised and received
  Received 8 messages, 3 notifications, 0 in queue
  Sent 13 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  IPV6_VRF_IBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 11:11:12::1, Local port: 179
Foreign host: 11:11:12::2, Foreign port: 49010
Nexthop: 7.7.7.7
Nexthop global: 11:11:12::1
Nexthop local: fe80::5054:ff:fe8b:8f5c
BGP connection: shared network
Last Reset: 00:02:16, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier.)
```

IPv6 EBGP Configuration

Below figure displays the minimum configuration required to enable BGP on an interface. R1, R2 and R3 are three routers belonging to the different AS, AS100 AS200 and AS300, connecting to network 11:11:11::/64 and 11:11:12::/64. First, define the routing process and the AS number to which the routers belong. Then, define BGP neighbors to start exchanging routing updates.

Topology

Figure 1-22: IPv6 EBGP peering

R1

`#configure terminal`	Enter Configuration mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ipv6 address 33::1/128`	Assign an IPv6 address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)#ipv6 address 11:11:11::1/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode.
`(config-router)#bgp router-id 1.1.1.1`	Assign a BGP router ID.
`(config-router)#neighbor 11:11:11::2 remote-as 200`	Configure BGP neighbor by specifying the neighbor IP address.
`(config-router)#address-family ipv6 unicast`	Enter the IPv6 Unicast Address Family.
`(config-router-af)#neighbor 11:11:11::2 activate`	Activate the neighbor in the address family.
`(config-router-af)#network 33::1/128`	Advertise the loopback network into BGP.
`(config-router-af)#end`	Exit from Router BGP mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ipv6 address 22::1/128`	Assign an IPv6 address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter interface mode for xe1
`(config-if)#ipv6 address 11:11:11::2/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode for xe2.
`(config-if)#ipv6 address 11:11:12::2/64`	Assign an IPv6 address to the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 200`	Enter Router BGP mode.
`(config-router)#bgp router-id 2.2.2.2`	Assign a BGP router ID.
`(config-router)#neighbor IPV6_EBGP_PEER peer-group range 11::1/16`	Configure a dynamic peer group, IPV6_EBGP_PEER.
`(config-router)#neighbor IPV6_EBGP_PEER remote-as 100`	Configure remote AS with peer group, IPV6_EBGP_PEER.
`(config-router)#neighbor IPV6_EBGP_PEER optional-as 300`	Configure optional AS with peer group, IPV6_EBGP_PEER.
`(config-router)#address-family ipv6 unicast`	Enter IPv6 Unicast Address Family.
`(config-router-af)#neighbor IPV6_EBGP_PEER activate`	Activate peer group in the address family.
`(config-router-af)#network 22::1/128`	Advertise the loopback network into BGP.
`(config-router-af)#end`	Exit from Router BGP mode.

R3

#Configure terminal	Enter Configuration mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ipv6 address 44::1/128	Assign an IPv6 address.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2
(config-if)#ipv6 address 11:11:12::3/64	Assign an IPv6 address.
(config-if)#exit	Exit interface mode.
(config)#router bgp 300	Enter Router BGP mode.
(config-router)#bgp router-id 3.3.3.3	Assign a BGP router ID.
(config-router)#neighbor 11:11:12::2 remote-as 200	Configure the BGP neighbor by specifying the neighbor IP address.
(config-router)#address-family ipv6 unicast	Enter the IPv6 Unicast Address Family.
(config-router-af)#neighbor 11:11:12::2 activate	Activate the neighbor in address family.
(config-router-af)#network 44::1/128	Advertise the loopback network into BGP.
(config-router-af)#end	Exit from Router BGP mode.

Validation

R2

```
#show ipv6 bgp sum
BGP router identifier 2.2.2.2, local AS number 200
BGP table version is 5
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
*11:11:11::1             4   100    9         11       5      0      0
00:01:28              1
*11:11:12::3             4   300    6          6       5      0      0
00:01:14              1
* Dynamically created based on a listen range command

BGP dynamic peer-group: IPV6_EBGP_PEER
  listen range:  11::/16
  Total number of dynamically created neighbors/limit: 2/(200)
Total number of dynamically created neighbors: 2
Total number of activated dynamic peer-groups for IPv6 Unicast address-family:
1
Total number of neighbors 2
Total number of Established sessions 2


#show running-config bgp
!
```

```
router bgp 200
 bgp router-id 2.2.2.2
 neighbor IPV6_EBGP_PEER peer-group range 11::/16
 neighbor IPV6_EBGP_PEER remote-as 100
 neighbor IPV6_EBGP_PEER optional-as 300
 !
 address-family ipv6 unicast
 network 22::1/128
 neighbor IPV6_EBGP_PEER activate
 exit-address-family
!

#show bgp ipv6 neighbors
BGP neighbor is 11:11:11::1, remote AS 100, local AS 200, external link
 Member of peer-group IPV6_EBGP_PEER for session parameters
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:02:15
  Last read 00:00:16, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 12 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 0, Offset 0, Mask 0x1
  IPV6_EBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 5, neighbor version 5
  Index 0, Offset 0, Mask 0x0
  IPV6_EBGP_PEER peer-group member
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 11:11:11::2, Local port: 53043
Foreign host: 11:11:11::1, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: 11:11:11::2
Nexthop local: fe80::5054:ff:fe95:85ec
BGP connection: shared network
Last Reset: 00:02:20, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)

BGP neighbor is 11:11:12::3, remote AS 300, local AS 200, external link
 Member of peer-group IPV6_EBGP_PEER for session parameters
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 00:02:01
  Last read 00:00:02, hold time is 90, keepalive interval is 30 seconds
```

```
 Neighbor capabilities:
   Route refresh: advertised and received (old and new)
   Address family IPv4 Unicast: advertised and received
   Address family IPv6 Unicast: advertised and received
 Received 8 messages, 0 notifications, 0 in queue
 Sent 8 messages, 0 notifications, 0 in queue
 Route refresh request: received 0, sent 0
 Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
 BGP table version 1, neighbor version 1
 Index 0, Offset 0, Mask 0x1
 IPV6_EBGP_PEER peer-group member
 Community attribute sent to this neighbor (both)
 0 accepted prefixes
 0 announced prefixes

For address family: IPv6 Unicast
 BGP table version 5, neighbor version 5
 Index 0, Offset 0, Mask 0x0
 IPV6_EBGP_PEER peer-group member
 Community attribute sent to this neighbor (both)
 1 accepted prefixes
 2 announced prefixes

 Connections established 1; dropped 0
Local host: 11:11:12::2, Local port: 47743
Foreign host: 11:11:12::3, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: 11:11:12::2
Nexthop local: fe80::5054:ff:fee5:b088
BGP connection: shared network
```

Enable eBGP Multihop

This example shows the minimum configuration required for enabling eBGP multihop on peers speaking BGP. eBGP multihop is used for routers that are not directly connected to each other. Typically, eBGP peers are directly connected, but if there is a requirement that necessitates this scenario, this configuration can be used.

Note: The IP addresses used in the configuration should be accessible through an IGP or static routing.

Topology

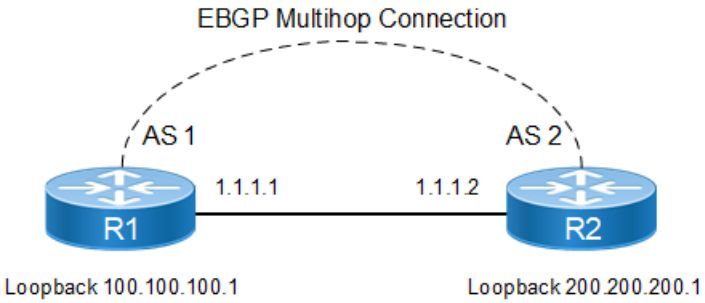

Figure 1-23: eBGP Multihop Connection

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 100.100.100.1/24 secondary`	Specify IP address to the interface.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#ip route 200.200.200.0/24 1.1.1.2`	Specify route IP address.
`(config)#router bgp 1`	Define the routing process. The number `1` specifies the AS number of `R1`.
`(config-router)#neighbor 200.200.200.1 remote-as 2`	Define BGP neighbors, and establish a TCP session. 200.200.200.1 is the IP address of the neighbor (R2), and 2 is the neighbor's AS number.
`(config-router)#neighbor 200.200.200.1 update-source lo`	Define BGP neighbors, to update the source routes.
`(config-router)#neighbor 200.200.200.1 ebgp-multihop`	Define the neighbor `200.200.200.1` for eBGP multihops.
`(config-router)#address-family ipv4 unicast`	Enter the Ipv4 Unicast Address Family.
`(config-router-af)# neighbor 200.200.200.1 activate`	Activate the neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address-family mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 200.200.200.1/24 secondary`	Specify IP address to the interface.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#ip route 100.100.100.0/24 1.1.1.1`	Specify route IP address.
`(config)#router bgp 2`	Define the routing process. The number 2 specifies the AS number of R1.

© 2025 IP Infusion Inc. Proprietary

(config-router)#neighbor 100.100.100.1 remote-as 1	Define BGP neighbors, and establish a TCP session. 100.100.100.1 is the IP address of the neighbor (R2), and 1 is the neighbor's AS number.
(config-router)#neighbor 100.100.100.1 update-source lo	Define BGP neighbors, to update the source routes.
(config-router)#neighbor 100.100.100.1 ebgp-multihop	Define the neighbor 100.100.100.1 for eBGP multihops.
(config-router)#address-family ipv4 unicast	Config redistribute under address-family
(config-router-af)# neighbor 100.100.100.1 activate	Activate the neighbor under address family mode
(config-router-af)#redistribute static	Redistribute static route
(config-router-af)#exit-address-family	Exit address-family mode

Validation

R1

```
#show ip bgp neighbors
BGP neighbor is 200.200.200.1, remote AS 2, local AS 1, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  BGP state = Established, up for 00:00:22
  Last read 00:00:22, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 2 messages, 0 notifications, 0 in queue
  Sent 3 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
  Update source is lo
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
  External BGP neighbor may be up to 255 hops away.
Local host: 100.100.100.1, Local port: 179
Foreign host: 200.200.200.1, Foreign port: 59458
Nexthop: 100.100.100.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


#show ip bgp
BGP table version is 4, local router ID is 192.168.52.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   100.100.100.0/24 200.200.200.1        0         100           0       2 ?

Total number of prefixes 1


#show ip bgp neighbors
BGP neighbor is 200.200.200.1, remote AS 2, local AS 1, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  BGP state = Established, up for 00:00:26
  Last read 00:00:26, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 5 messages, 0 notifications, 0 in queue
  Sent 6 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
  Update source is lo
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
  External BGP neighbor may be up to 255 hops away.
Local host: 100.100.100.1, Local port: 57260
Foreign host: 200.200.200.1, Foreign port: 179
Nexthop: 100.100.100.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:31, due to BGP Notification sent
Notification Error Message: (Cease/Administratively Reset.)
```

R2

```
#sh ip bgp neighbors
BGP neighbor is 100.100.100.1, remote AS 1, local AS 2, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  BGP state = Established, up for 00:00:35
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
```

```
  Received 6 messages, 1 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
  Update source is lo
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
  External BGP neighbor may be up to 255 hops away.
Local host: 200.200.200.1, Local port: 179
Foreign host: 100.100.100.1, Foreign port: 57260
Nexthop: 200.200.200.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:40, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)

#show ip bgp
BGP table version is 4, local router ID is 192.168.52.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   100.100.100.0/24 1.1.1.1              0         100           32768    ?

Total number of prefixes 1
```

Enable Peer Groups

A BGP speaker might have the same update policies for a set of its peers. This is very useful if you have to change the update policies for all of the peers: Changing individual routers for separate policies can be very time-consuming, thus, peer groups play an important role in creating and assigning policies to a group of routers.

Note: A maximum of 32 peers can be bound in a peer group.

The peer group can be created dynamically or statically.

For dynamic peer groups all configuration can be done at group level only.

The static peer group configuration falls into two categories:

- Attributes that can be configured only at group level. Attempt to configure at member peer level will return error.
- Attributes that allow member peer level configuration. The member peer configuration has precedence.

Category A: Neighbor configuration allowed only at peer-group level

The outbound update impacting configurations of peer group will replace peer member configurations of same attributes when a peer becomes member of peer group. Outbound attribute modifications to group members are not allowed.

Following are the commands which are allowed at peer-group level.

- neighbor WORD activate - neighbor activate
- neighbor WORD as-origination-interval <1-65535> - neighbor as-origination-interval
- neighbor WORD attribute-unchanged ({as-path|next-hop|med}|) - neighbor attribute-unchanged
- neighbor WORD fall-over bfd - neighbor fall-over bfd
- neighbor WORD fall-over bfd multihop - neighbor enforce-multihop
- neighbor WORD next-hop-self - neighbor next-hop-self
- neighbor WORD remove-private-AS - neighbor remove-private-AS
- neighbor WORD route-reflector-client - neighbor route-reflector-client
- neighbor WORD route-server-client - neighbor route-server-client
- neighbor WORD send-community - neighbor send-community
- neighbor WORD distribute-list WORD out - neighbor distribute-list
- neighbor WORD dont-capability-negotiate - neighbor dont-capability-negotiate
- neighbor WORD capability orf prefix-list (both|receive|send) - neighbor capability orf prefix-list
- neighbor WORD filter-list WORD out - neighbor filter-list
- neighbor WORD prefix-list WORD out - neighbor prefix-list
- neighbor WORD route-map WORD out - neighbor restart-time
- neighbor WORD advertisement-interval <1-65535> - neighbor advertisement-interval
- neighbor WORD disallow-infinite-holdtime - neighbor disallow-infinite-holdtime
- neighbor WORD local-as <1-4294967295> - neighbor local-as

Category: Neighbor configuration allowed at peer-group member level; precedence based

For the below configuration, member level configurations will take precedence over peer group configuration.

Following are the commands which are allowed at member level also.

- neighbor WORD authentication-key WORD - neighbor authentication-key
- neighbor WORD remote-as <1-4294967295> - neighbor remote-as
- neighbor WORD allowas-in <1-10> - neighbor allowas-in
- neighbor WORD description WORD - neighbor description
- neighbor WORD distribute-list WORD In - neighbor distribute-list
- neighbor WORD ebgp-multihop - neighbor ebgp-multihop
- neighbor WORD ebgp-multihop <1-255> - neighbor ebgp-multihop
- neighbor WORD maximum-prefix <1-4294967295> - neighbor maximum-prefix
- neighbor WORD update-source WORD - neighbor update-source
- neighbor WORD weight <0-65535> - neighbor weight
- neighbor WORD shutdown - neighbor shutdown

- neighbor WORD strict-capability-match - neighbor strict-capability-match
- neighbor WORD route-map WORD in - neighbor restart-time
- neighbor WORD prefix-list WORD in - neighbor prefix-list
- neighbor WORD passive - neighbor passive
- neighbor WORD override-capability - neighbor override-capability
- neighbor WORD filter-list WORD in - neighbor filter-list
- neighbor WORD enforce-multihop - neighbor enforce-multihop
- neighbor WORD collide-established - neighbor collide-established

BGP Peer Groups for Address-Family IPv4 Unicast

In the following scenario, R1, R2, and R3 belong to the same peer group ABC. R1, R2 and R3 are in AS 200 and R1 is the route reflector.R4 and R1 are eBGP peers. R4 is in AS 100.

Topology

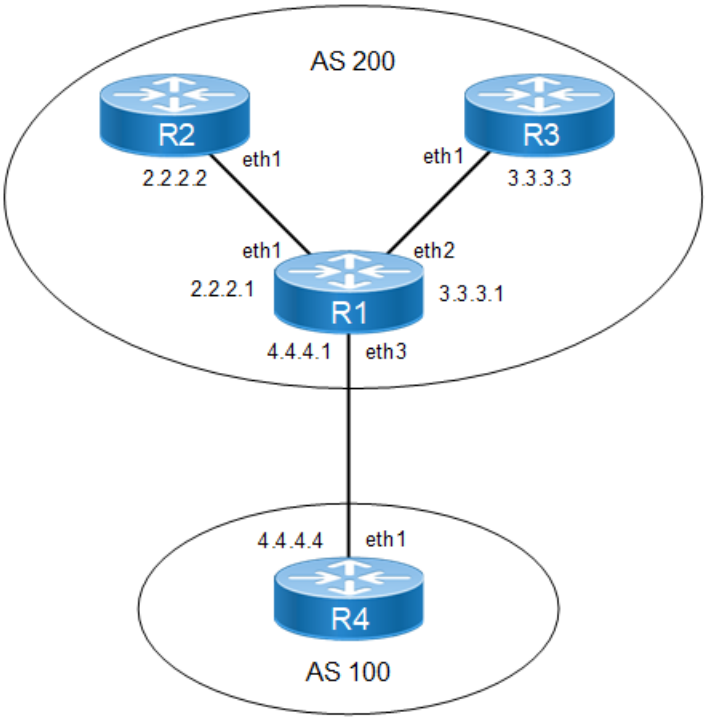

Figure 1-24: BGP Peer Groups with IPv4 Unicast Members

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number 200 specifies the AS number of R2.
`OcNOS(config-router)#neighbor ABC peer-group`	Configuring ABC peer-group
`(config-router)#neighbor ABC remote-as 200`	Assign options to the peer group named ABC.

(config-router)#neighbor 2.2.2.2 peer-group ABC	Define neighbor 2.2.2.2 (R2) as a peer group member.
(config-router)#neighbor 3.3.3.3 peer-group ABC	Define neighbor 3.3.3.3 (R3) as a peer group member.
(config-router)#neighbor 4.4.4.4 remote-as 100	Define neighbor 4.4.4.4 (R4) is the IP address of R4 and 100 is the AS number.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor ABC activate	Activate neighbor under address family mode
(config-router-af)# neighbor 4.4.4.4 activate	Activate neighbor under address family mode
(config-router-af)# neighbor ABC route-reflector-client	Configure the peer-group ABC to be route-reflector-client
(config-router-af)#network 1.1.1.1/32	Advertise the network 1.1.1.1/32
(config-router-af)#network 11.11.11.11/32	Advertise the network 11.11.11.11/32

R2

#configure terminal	Enter configure mode.
(config)#router bgp 200	Define the routing process. The number 200 specifies the AS number of R2.
(config-router)#neighbor 2.2.2.1 remote-as 200	Create a TCP connection with neighbor 2.2.2.1 of AS 200.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 2.2.2.1 activate	Activate neighbor under address family mode
OcNOS(config-router-af)#exit-address-family	Exit address family mode

R3

#configure terminal	Enter configure mode.
(config)#router bgp 200	Define the routing process. The number 200 specifies the AS number of R3.
(config-router)#neighbor 3.3.3.1 remote-as 200	Create a TCP connection with neighbor 3.3.3.1 of AS 200.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 3.3.3.1 activate	Activate neighbor under address family mode
OcNOS(config-router-af)#exit-address-family	Exit address family mode

R4

#configure terminal	Enter configure mode.
(config)#router bgp 100	Define the routing process. The number 100 specifies the AS number of R4.
(config-router)#neighbor 4.4.4.1 remote-as 200	Create a TCP connection with neighbor 4.4.4.1 of AS 200.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode

`(config-router-af)# neighbor 4.4.4.1 activate`	Activate neighbor under address family mode
`(config-router-af)#exit-address-family`	Exit address family mode

Validation

R1

```
R1#show ip bgp neighbors
BGP neighbor is 2.2.2.2, remote AS 200, local AS 200, internal link
 Member of peer-group ABC for session parameters
  BGP version 4, local router ID 192.168.52.2, remote router ID 10.12.7.155
  BGP state = Established, up for 00:04:55
  Last read 00:04:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes
 Connections established 1; dropped 0
Local host: 2.2.2.1, Local port: 33865
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 2.2.2.1
Nexthop global: 1111::1
Nexthop local: fe80::a00:27ff:fecc:47a6
BGP connection: non shared network

BGP neighbor is 3.3.3.3, remote AS 200, local AS 200, internal link
 Member of peer-group ABC for session parameters
  BGP version 4, local router ID 192.168.52.2, remote router ID 10.12.7.153
  BGP state = Established, up for 00:04:55
  Last read 00:04:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
```

```
  0 accepted prefixes
  2 announced prefixes
 Connections established 1; dropped 0
Local host: 3.3.3.1, Local port: 44280
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 3.3.3.1
Nexthop global: fe80::a00:27ff:fe85:25d4
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 4.4.4.4, remote AS 100, local AS 200, external link
  BGP version 4, remote router ID 10.12.7.120
  BGP state = Established, up for 00:04:55
  Last read 00:04:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes
 Connections established 1; dropped 0
Local host: 4.4.4.1, Local port: 55493
Foreign host: 4.4.4.4, Foreign port: 179 Nexthop: 4.4.4.1
Nexthop global: fe80::a00:27ff:fe7e:674a
Nexthop local: ::
BGP connection: non shared network

R1#show ip bgp summary
BGP router identifier 192.168.52.2, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries
Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
2.2.2.2                   4   200   12         12       1      0      0
00:05:02               0
3.3.3.3                   4   200   12         12       1      0      0
00:05:02               0
4.4.4.4                   4   100   12         12       1      0      0
00:05:02               0
Total number of neighbors 3
Total number of Established sessions 3
```

R2

```
R2#show ip bgp
BGP table version is 4, local router ID is 10.12.65.123
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  1.1.1.1/32       2.2.2.1               0         100           0        i
*>i  11.11.11.11/32   2.2.2.1               0         100           0        i

Total number of prefixes 2
R2#
```

R3

```
R3#show ip bgp
BGP table version is 8, local router ID is 10.12.65.121
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  1.1.1.1/32       3.3.3.1               0         100           0        i
*>i  11.11.11.11/32   3.3.3.1               0         100           0        i

Total number of prefixes 2
R3#
```

Peer-group can have either iBGP or eBGP peers but not both.

Validation

The configuration above fails with an appropriate error:

```
R1(config)#router bgp 200
R1(config-router)#neighbor 4.4.4.4 peer-group ABC
%% Peer with AS 100 cannot be in this peer-group, members must be all internal
or all external
R1(config-router)#
```

Peer group members inherit the properties of Outbound Policies configured for Peer-group.

R1

`#configure terminal`	Enter configure mode.
`(config)# ip access-list permit-1`	Configure access-list to permit 1.1.1.1/32
`(config-ip-acl)# permit any 1.1.1.1/32 any`	Configure a permit statement in the acl to permit 1.1.1.1/32
`(config-ip-acl)#exit`	Exit ip access-list mode
`(config)# route-map permit-only-1`	Configure route-map
`(config-route-map)# match ip address permit-1`	Configure a match statement in the route-map to match the access-list permit-1
`(config-route-map)#set local-preference 250`	Set local preference as 250
`(config-route-map)#exit`	Exit route-map mode
`(config)#router bgp 200`	Define the routing process. The number 200 specifies the AS number of R1.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode

(config-router-af)# neighbor ABC route-map permit-only-1 out	Configure the peer-group ABC with route-map in the outbound direction
(config-router-af)#exit	Exit router BGP mode
(config-router)#exit	Exit router mode.
(config# exit	Exit configure terminal mode
#clear ip bgp peer-group ABC soft out	Do outbound soft reset for the peer-group ABC for the policy to take affect for the peer-group members

Validation

R1

```
R1#show bgp neighbors 2.2.2.2
BGP neighbor is 2.2.2.2, remote AS 200, local AS 200, internal link
 Member of peer-group ABC for session parameters
  BGP version 4, local router ID 10.12.65.126, remote router ID 10.12.65.123
  BGP state = Established, up for 00:07:01
  Last read 00:00:15, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 20 messages, 0 notifications, 0 in queue
  Sent 28 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Outbound path policy configured
  Route map for outgoing advertisements is *permit-only-1
  0 accepted prefixes
  1 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.1, Local port: 179
Foreign host: 2.2.2.2, Foreign port: 42657
Nexthop: 2.2.2.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:08:39, due to Hold Timer Expired (Notification sent)
Notification Error Message: (Hold Timer Expired/No sub-error code)

R1#show bgp neighbors 3.3.3.3
BGP neighbor is 3.3.3.3, remote AS 200, local AS 200, internal link
 Member of peer-group ABC for session parameters
  BGP version 4, local router ID 10.12.65.126, remote router ID 10.12.65.121
  BGP state = Established, up for 00:11:46
  Last read 00:00:18, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
```

```
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 29 messages, 0 notifications, 0 in queue
  Sent 32 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Outbound path policy configured
  Route map for outgoing advertisements is *permit-only-1
  0 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 3.3.3.1, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 48008
Nexthop: 3.3.3.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

R2

```
R2#show ip bgp
BGP table version is 3, local router ID is 10.12.65.123
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  1.1.1.1/32       2.2.2.1               0         250           0        i

Total number of prefixes 1
```

R3

```
R3#show ip bgp
BGP table version is 7, local router ID is 10.12.65.121
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>i  1.1.1.1/32       3.3.3.1               0         250           0        i

Total number of prefixes 1
```

Peer-group-members cannot be configured with Outbound Policies.

Validation

The above configuration fails with an appropriate error:

```
R1(config)#router bgp 200
R1(config)#address-family ipv4 unicast
R1(config-router-af)#neighbor 2.2.2.2 route-map permit-only-11 out
%% Invalid command for a peer-group member
```

Peer-group-members inherit the properties of Inbound Policies configured for Peer-group.

R2

`#configure terminal`	Enter configure mode.
`(config)# interface lo`	Enter interface mode for Loopback interface
`(config-if)#ip address 100.1.1.1/24 secondary`	Configure IP address for Loopback interaface
`(config-if)#interface eth3`	Enter interface mode for interface eth3
`(config-if)#ip address 22.1.1.1/24`	Configure IP address for interface eth3
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Enter router bgp mode
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#network 22.1.1.0/24`	Advertise the network of eth3 in BGP
`(config-router-af)#network 100.1.1.0/24`	Advertise the network of Loopback in BGP
`(config-router-af)#exit-address-family`	Exit router bgp mode

R3

`#configure terminal`	Enter configure mode.
`(config)# interface lo`	Enter interface mode for Loopback interface
`(config-if)#ip address 100.1.1.2/24 secondary`	Configure IP address for Loopback interaface
`(config-if)#interface eth3`	Enter interface mode for interface eth3
`(config-if)#ip address 22.1.1.2/24`	Configure IP address for interface eth3
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Enter router bgp mode
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#network 22.1.1.0/24`	Advertise the network of eth3 in BGP
`(config-router-af)#network 100.1.1.0/24`	Advertise the network of Loopback in BGP
`(config-router-af)#exit`	Exit router bgp mode

R1

`#configure terminal`	Enter configure mode.
`(config)#ip access-list permit-22`	Configure access-list to permit 22.1.1.0/24
`(config-ip-acl)# permit any 22.1.1.0/24 any`	Configure a permit statement to permit 22.1.1.0/24
`(config-ip-acl)#exit`	Exit ip access-list mode
`(config)#route-map permit-only-22`	Configure route-map

(config-route-map)#match ip address permit-22	Configure match statement in route-map to match the access-list permit-22
(config-route-map)#exit	Exit route-map mode
(config)#router bgp 200	Enter BGP router mode
(config)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor ABC activate	Activate peer-group in the address family
(config-router-af)#neighbor ABC route-map permit-only-22 in	Configure the peer-group ABC with route-map in the inbound direction
(config-router)#exit	Exit router bgp mode
(config)# exit	Exit configure terminal mode
#clear ip bgp peer-group ABC soft in	Do inbound soft reset for the peer-group ABC for the policy to take affect for the peer-group members

Validation

R1

```
R1#show ip bgp

BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, g group-best, *
valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network              Next Hop              Metric        LocPrf
Weight  Path
*>   1.1.1.1/32          0.0.0.0               0             100          32768
i
*>   11.11.11.11/32      0.0.0.0               0             100          32768
i
*>i  22.1.1.0/24         2.2.2.2               0             100          0
i
* i                       3.3.3.3               0             100          0
i

Total number of prefixes 3
```

Peer group members can be configured with Inbound Policies

R1

#configure terminal	Enter configure mode.
(config)#ip access-list permit-100	Configure access-list to permit 100.1.1.0/24
(config-ip-acl)# permit any 100.1.1.0/24 any	Configure a permit statement to permit 100.1.1.0/24
(config-ip-acl)#exit	Exit ip access-list mode
(config)#route-map permit-only-100	Configure route-map
(config-route-map)#match ip address permit-100	Configure match statement in route-map to match the access-list permit-100

(config-route-map)#exit	Exit route-map mode
(config)#router bgp 200	Enter BGP router mode
(config)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 2.2.2.2 activate	Activate neighbor in the address family
(config-router-af)#neighbor 2.2.2.2 route-map permit-only-100 in	Configure the peer-group-member R2(2.2.2.2) with route-map in the inbound direction
(config-router-af)#exit	Exit address-family mode
(config-router)#exit	Exit router bgp mode
(config)# exit	Exit configure terminal mode
#clear ip bgp peer-group ABC soft in	Do inbound soft reset for the peer-group ABC for the policy to take affect for the peer-group members

Validation

R1

```
R1#show ip bgp
BGP table version is 4, local router ID is 10.12.65.126
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   1.1.1.1/32       0.0.0.0               0         100            32768    i
*>   11.11.11.11/32   0.0.0.0               0         100            32768    i
*>i  22.1.1.0/24      3.3.3.3               0         100            0        i
*>i  100.1.1.0/24     2.2.2.2               0         100            0        i

Total number of prefixes 4
R1#
```

Route Redistribution in BGP

If there are routers that run both OSPF and BGP, certain OSPF routes might have to be sent to other eBGP peers. This can be achieved using the redistribution feature. Consider the following topology, in which R1 and R2 are eBGP peers, and R2 and R3 are OSPF peers. R2 is redistributing OSPF routes into BGP. The OSPF routes are sent to the R1 BGP routing table. This configuration assumes that all OSPF and eBGP sessions are up and running, and that only the redistribution must be configured.

Topology

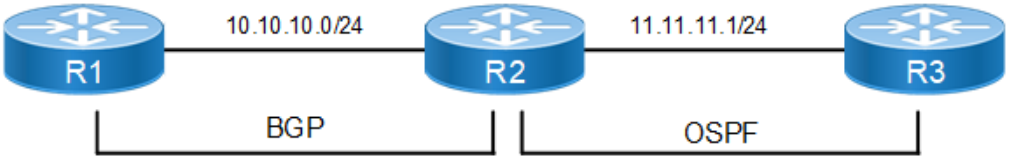

Figure 1-25: Redistribute with OSPF

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 100`	Define the routing process. The number `100` specifies the AS number of R2.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#redistribute ospf`	Redistribute OSPF routes in the R2 routing table into the R1 BGP routing table.
`(config-router-af)# exit-address-family`	Exit address-family mode.

Validation

```
#show ip bgp
BGP table version is 3, local router ID is 192.168.52.2
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   3.3.3.3/32       10.10.10.2            11        100            0       100
?
*>   11.11.11.0/24    10.10.10.2            1         100            0       100
?

Total number of prefixes 2
```

Add Multiple Instances of the Same Autonomous System

BGP supports adding the same AS number multiple times to influence the route selection process. This can be done using route maps, as described below.

Under normal circumstances, any route advertised by R1 is sent to R4 via two different routes, and then R4 selects the path from R2. This decision can be influenced by adding multiple instances of AS number 200 at R2.

Topology

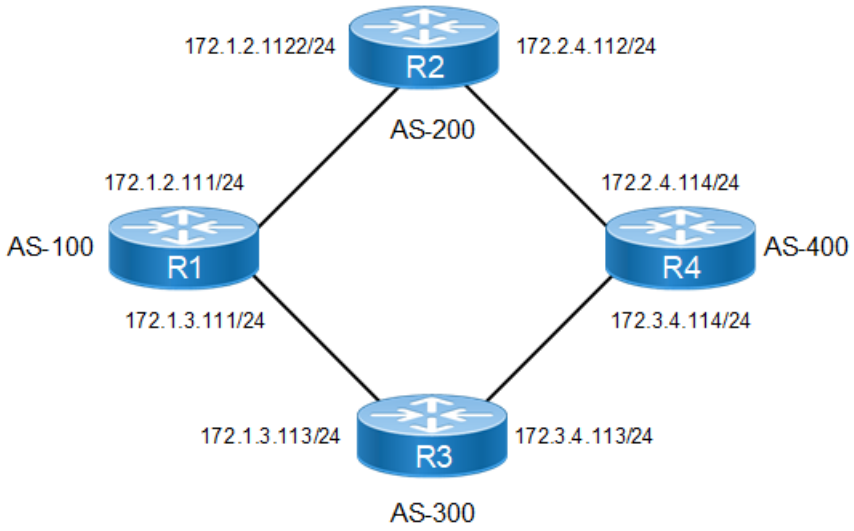

Figure 1-26: Multiple Instances of Same AS

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 44.44.44.1/24 secondary`	Specify the IP address for the interface.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#router bgp 100`	Define the routing process with AS number 100.
`(config-router)#neighbor 172.1.2.112 remote-as 200`	Define neighbor R2. 172.1.2.112 is the IP address of R2, and 200 is the AS number.
`(config-router)#neighbor 172.1.3.113 remote-as 300`	Define neighbor R3. 172.1.3.113 is the IP address of R2, and 300 is the AS number.
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.1.2.112 activate`	Activate neighbor under address family mode
`(config-router-af)# neighbor 172.1.2.113 activate`	Activate neighbor under address family mode
`(config-router-af)#network 44.44.44.0/24`	Advertise network 44.44.44.0/24 through BGP. This route reaches R4 via R2 and R3.

R2

`#configure terminal`	Enter configure mode.
`(config)#route-map mul_inst permit 10`	Define the route-map multiple instance with permit definition sequence number 10.

(config-route-map)#set as-path prepend 200 200	Prepend AS number 200 two times to the AS_PATH attribute in the BGP Update message.
(config-route-map)#exit	Exit Route-map mode, and return to Configure mode.
(config)#router bgp 200	Define the routing process with AS number 200.
(config-router)#neighbor 172.1.2.111 remote-as 100	Define neighbor R1. 172.1.2.111 is the IP address of R1, and 100 is the AS number.
(config-router)#neighbor 172.2.4.114 remote-as 400	Define neighbor R4. 172.2.4.114 is the IP address of R2, and 400 is the AS number.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 172.2.4.114 activate	Activate neighbor under address family mode
(config-router-af)# neighbor 172.1.2.111 activate	Activate neighbor under address family mode
(config-router-af)#neighbor 172.2.4.114 route-map mul_inst out	Apply route-map multi_inst to all outbound routes to R4

R3

#configure terminal	Enter configure mode.
(config)#router bgp 300	Define the routing process with AS number 300.
(config-router)#neighbor 172.1.3.111 remote-as 100	Define neighbor R1. 172.1.3.111 is the IP address of R1, and 100 is the AS number.
(config-router)#neighbor 172.3.4.114 remote-as 400	Define neighbor R4. 172.3.4.114 is the IP address of R4, and 400 is the AS number.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 172.3.4.114 activate	Activate neighbor under address family mode
(config-router-af)# neighbor 172.1.3.111 activate	Activate neighbor under address family mode

R4

#configure terminal	Enter configure mode.
(config)#router bgp 400	Define the routing process with AS number 400.
(config-router)#neighbor 172.2.4.112 remote-as 200	Define neighbor R2. 172.2.4.112 is the IP address of R2, and 200 is the AS number.
(config-router)#neighbor 172.3.4.113 remote-as 300	Define neighbor R3. 172.3.4.113 is the IP address of R3, and 300 is the AS number.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 172.2.4.112 activate	Activate neighbor under address family mode
(config-router-af)# neighbor 172.3.4.113 activate	Activate neighbor under address family mode

Validation

```
#show ip bgp
BGP table version is 1, local router ID is 44.44.44.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*> 44.44.44.0/24 0.0.0.0 0 100 32768 i

Total number of prefixes 1
```

Remove the Multi-Exit Disc Attribute from Update Messages

You can remove the Multi-Exit Disc (MED) attribute values from received update messages.

Topology

Figure 1-27: Remove MED Attribute

R1

`#configure terminal`	Enter configure mode.
`(config)#route-map med permit 1`	Define the route-map MED with permit definition sequence number 1.
`(config-route-map)#set metric 400`	Set the metric value.
`(config-route-map)#exit`	Exit Route-map mode, and return to Configure mode.
`(config)#router bgp 100`	Define the routing process with AS number 100.
`(config-router)#neighbor 1.1.1.2 remote-as 200`	Define neighbor R2. 1.1.1.2 is the IP address of R2, and 200 is the AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 1.1.1.2 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode
`(config-router)#exit`	Exit router BGP mode
`(config)#interface xe2`	Enter interface mode
`(config-if)#ip ad 10.10.10.1/24`	Assign IP address
`(config-if)#no shutdown`	Make interface administratively up
`(config-if)#exit`	Exit interface mode
`(config)#ip route 100.0.0.0/8 10.10.10.2`	Configure the static route with the nexthop address.

R3

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process with AS number 200.
`(config-router)#neighbor 2.2.2.1 remote-as 200`	Define neighbor R2. 2.2.2.1 is the IP address of R2, and 200 is the AS number.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 2.2.2.1 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

Removing Sent and Received MED values

The following describes how to remove the received and sent MED values, respectively.

R2 - Remove Received MED Value

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process with AS number 200.
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Define neighbor R1. 1.1.1.1 is the IP address of R1, and 100 is the AS number.
`(config-router)#neighbor 2.2.2.2 remote-as 200`	Define neighbor R3. 2.2.2.2 is the IP address of R3, and 200 is the AS number.
`(config-router)#bgp bestpath med remove-recv- med`	Enable the remove received MED value option.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 1.1.1.1 activate`	Activate neighbor under address family mode
`(config-router-af)# neighbor 2.2.2.2 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

R1 - Add Static Route

`(config)#router bgp 100`	Enter to router bgp mode
`(config)#address-family ipv4 unicast`	Config redistribute under address-family
`(config-router-af)#redistribute static`	Redistribute the static routes

R2 - Remove Send MED Value

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process with AS number 200.
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Define neighbor R1. 1.1.1.1 is the IP address of R1, and 100 is the AS number.

`(config-router)#neighbor 2.2.2.2 remote-as 200`	Define neighbor R3. 2.2.2.2 is the IP address of R3, and 200 is the AS number.
`(config-router)#bgp bestpath med remove-send- med`	Enable the remove sent MED value option.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 1.1.1.1 activate`	Activate neighbor under address family mode
`(config-router-af)# neighbor 2.2.2.2 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

Validation

```
R2#show ip bgp
BGP table version is 2, local router ID is 192.168.52.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   100.0.0.0        1.1.1.1         removed     100           0       100 ?

Total number of prefixes 1


R3#show ip bgp
BGP table version is 1, local router ID is 192.168.52.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
* i  100.0.0.0        1.1.1.1               400       100           0       100 ?

Total number of prefixes 1
```

BGP Four-Byte Autonomous System

Extended AS numbers can be mapped to 2-byte AS numbers if the value is less than, or equal to, 65535. If the AS number is higher than 65535, it cannot be mapped to a 2-byte AS number. Therefore, if a BGP speaker is configured with a non-mappable AS number, it must enable the BGP extended ASN capability in OcNOS.

Note: Autonomous System number 23456 is a reserved IANA number for AS transition; thus, it is recommended that no system be configured with 23456 as its AS number.

The extended ASN capability is disabled by default. However, when it is enabled, it is able to interoperate with a 2-byte AS-numbered speaker, in compliance with RFC 4893.

If a 4-byte AS number is configured in the provider's network using the standard IPv4/IPv6 BGP, it is recommended that the PE routers be 4-byte AS-enabled before connecting to 4-byte AS-enabled customer networks. For implications related to AS number transition issues, refer to RFC 4893.

You can also set up 4-byte AS-specific extended communities and route distinguishers (RDs) with limited capabilities. However, it is recommended that 2-byte AS-specific RDs and extended communities be used for regular deployment.

BGP encodes an ASN into four octets, so that more autonomous systems can be supported. Extended ASN capability is advertised in the Open message capabilities when the 4-octet ASN capability is enabled. When the 4-octet ASN capability is enabled, the valid ASN value range is <1-4294967295>, with the exception discussed in the first Note, above.

Note: Four-octet capability is disabled by default.

4-Octet ASN Capability Enabled on R1 and R2

In this example, 4-Octet ASN capability is enabled on BGP speakers R1 and R2.

Topology

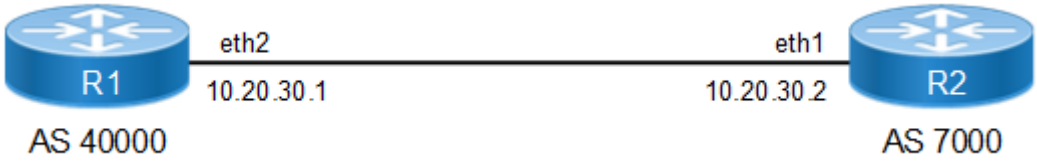

Figure 1-28: 4-Octet ASN on Both Routers

R1

`#configure terminal`	Enter configure mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability.
`(config)#router bgp 400000`	Assign the ASN value (400000) to the router.
`(config-router)#neighbor 10.20.30.2 remote-as 7000`	Specify the neighbor's IP address (10.20.30.2) and the ASN value of the neighbor (7000).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 10.20.30.2 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

R2

`#configure terminal`	Enter configure mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability.
`(config)#router bgp 7000`	Assign the ASN value (7000) to the router.
`(config-router)#neighbor 10.20.30.1 remote-as 400000`	Specify the neighbor's IP address (10.20.30.1) and the ASN value of the neighbor (400000).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode

`(config-router-af)# neighbor 10.20.30.1 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

4-Octet ASN Capability Enabled on R1 and Disabled on R2

In the following two examples, 4-Octet ASN capability is enabled on BGP speaker R1 and disabled on R2.

Topology

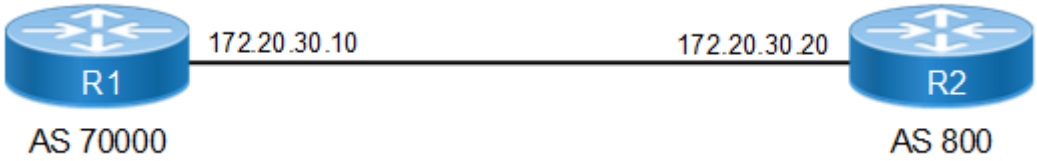

Figure 1-29: 4-Octet ASN on One Router

R1

`#configure terminal`	Enter configure mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability.
`(config)#router bgp 70000`	Assign the ASN value (70000) to the router.
`(config-router)#neighbor 172.20.30.20 remote-as 800`	Specify the neighbor's IP address (172.20.30.20) and the ASN value of the neighbor (800).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.20.30.20 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

R2

`#configure terminal`	Enter configure mode.
`(config)#no bgp extended-asn-cap`	Disable 4-octet ASN capability.
`(config)#router bgp 800`	Assign the ASN value (800) to the router.
`(config-router)#neighbor 172.20.30.10 remote-as 70000`	Specify the neighbor's IP address (172.20.30.10) and the ASN value of the neighbor (70000).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.20.30.10 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

© 2025 IP Infusion Inc. Proprietary

Topology

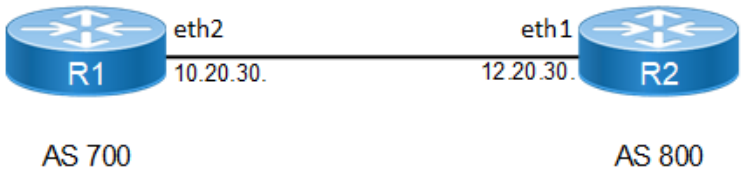

Figure 1-30: 4-Octet ASN

R1

`#configure terminal`	Enter configure mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability.
`(config)#router bgp 700`	Assign the ASN value (700) to the router.
`(config-router)#neighbor 172.20.30.20 remote-as 800`	Specify the neighbor's IP address (172.20.30.20) and the ASN value of the neighbor (800).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.20.30.20 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

R2

`#configure terminal`	Enter configure mode.
`(config)#no bgp extended-asn-cap`	Disable 4-octet ASN capability.
`(config)#router bgp 800`	Assign the ASN value (800) to the router.
`(config-router)#neighbor 172.20.30.10 remote-as 700`	Specify the neighbor's IP address (172.20.30.10) and the ASN value of the neighbor (700).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.20.30.10 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

Validation

```
#show ip bgp summary
BGP router identifier 192.168.52.2, local AS number 400000
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS    MsgRcv    MsgSen TblVer    InQ    OutQ    Up/
Down    State/PfxRcd
10.20.30.2               4  7000     2          3       1       0       0
00:00:08               0

Total number of neighbors 1

Total number of Established sessions 1
```

```
#show ip bgp neighbors
BGP neighbor is 10.20.30.1, remote AS 400000, local AS 7000, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  BGP state = Established, up for 00:02:20
  Last read 00:00:20, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 6 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.20.30.2, Local port: 49434
Foreign host: 10.20.30.1, Foreign port: 179
Nexthop: 10.20.30.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Extended Community with a 4-Byte ASN

In the following example, CE1, PE1, PE2, and CE2 support 4-byte ASN capability.

Note: PE1 and PE2 should both either be 4-byte-ASN capable or 2-byte-ASN capable. Support for the combination of one 4-byte-ASN capable PE with one 2-byte-ASN-capable PE is currently unavailable.

Topology

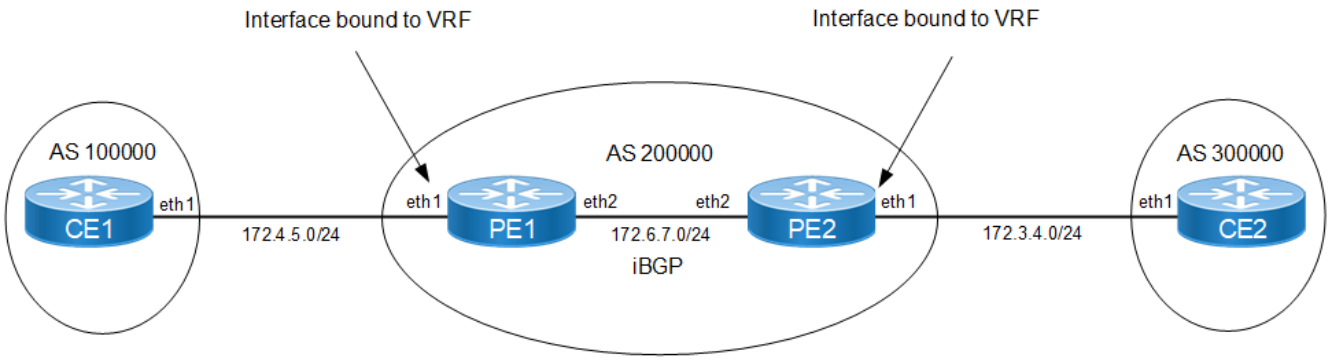

Figure 1-31: Extended Communities — 4-Byte ASN

CE1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 172.4.5.115/24`	Configure the IP address on this interface
`(config-if)#exit`	Exit interface mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability.
`(config)#router bgp 100000`	Assign the ASN value (100000) to the router. The ASN range is <1-4294967295>.
`(config-router)#neighbor 172.4.5.116 remote-as 200000`	Specify the neighbor's IP address (172.4.5.116) and the ASN value of the neighbor (200000).
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 172.4.5.116 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode

PE1

`#configure terminal`	Enter configure mode.
`(config)#bgp extended-asn-cap`	Enable 4-octet ASN capability. Dynamic change from 2- byte to 4-byte capability, or vice versa, is not allowed, unless the VRF is removed.
`(config)#ip vrf VRF1`	Specify the name of the VRF (VRF1) to be created.
`(config-vrf)#rd 1.1.1.1:200`	Assign a 4-byte route distinguisher (RD) for the VRF, which is a unique value on the router. The RD value can be in A.B.C.D:NN format.

(config-vrf)#route-target both 1.1.1.1:200	Specify the 4-Octet AS specific or IPv4 specific Transitive Route-Target extended community attribute.
(config-vrf)#exit	Exit VRF mode, and return to Configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip vrf forwarding VRF1	Bind the interface (eth1) to the VRF (VRF1).
(config-if)#ip address 172.4.5.116/24	Configure the IP address on this interface
(config-if)#exit	Exit interface mode.
(config)#ip route vrf VRF1 50.1.1.0/24 eth1	Create a VRF static route.
(config)#interface eth2	Enter interface mode.
(config-if)#ip address 172.6.7.116/24	Configure the IP address on this interface
(config-if)#exit	Exit interface mode.
(config)#router bgp 200000	Assign the ASN value (200000) to the router.
(config-router)#neighbor 172.6.7.117 remote-as 200000	Specify the neighbor's (PE2) IP address (172.6.7.117) and the ASN value of the neighbor (200000). In this case, it is an iBGP connection, so both PE1 and PE2 are in the same AS.
(config-router)#address-family vpnv4 unicast	Enable the exchange of VPNv4 routing information among ISP PE-routers, and enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 172.6.7.117 activate	Enable the exchange of routing information with a peer router.
(config-router-af)#exit	Exit Address-Family-VPNv4 mode.
(config-router)#address-family ipv4 vrf VRF1	Enable the exchange of VRF routing information among ISP PE-routers, and enter Address-Family-VRF mode.
(config-router-af)#neighbor 172.4.5.115	
remote-as 100000	Specify the neighbor's (CE1) IP address and ASN value.
(config-router-af)# neighbor 172.4.5.115 activate	Activate neighbor in address family mode
(config-router-af)#neighbor 172.4.5.115 send- community both	Enable extended community attribute for the neighbor.
(config-router-af)#redistribute static	Configure static redistribution.
(config-router-af)#end	Exit Address-Family mode

PE2

#configure terminal	Enter configure mode.
(config)#bgp extended-asn-cap	Enable 4-octet ASN capability. Dynamic change from 2- byte to 4-byte capability, or vice versa, is not allowed, unless the VRF is removed.
(config)#ip vrf VRF1	Specify the name of the VRF (VRF1) to be created.
(config-vrf)#rd 1.1.1.1:200	Assign a 4-byte route distinguisher (RD) for the VRF.
(config-vrf)#route-target both 1.1.1.1:200	Specify the 4-Octet AS specific or IPv4 specific Transitive Route-Target extended community attribute.
(config-vrf)#exit	Exit VRF mode, and return to Configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip vrf forwarding VRF1	Bind the interface (eth1) to the VRF (VRF1).
(config-if)#ip address 172.3.4.117/24	Configure the IP address on this interface

(config-if)#exit	Exit interface mode.
(config)#ip route vrf VRF1 200.1.1.0/24 eth1	Create a VRF static route.
(config)#interface eth2	Enter interface mode.
(config-if)#ip address 172.6.7.117/24	Configure the IP address on this interface
(config-if)#exit	Exit interface mode.
(config)#router bgp 200000	Assign the ASN value (200000) to the router.
(config-router)#neighbor 172.6.7.116 remote-as 200000	Specify the neighbor's (PE2) IP address (172.6.7.116) and the ASN value of the neighbor (200000). In this case, it is an iBGP connection, so both PE1 and PE2 are in the same AS.
(config-router)#address-family vpnv4 unicast	Enable the exchange of VPNv4 routing information among ISP PE-routers, and enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 172.6.7.116 activate	Enable the exchange of routing information with a peer router.
(config-router-af)#exit	Exit Address-Family-VPNv4 mode.
(config-router)#address-family ipv4 vrf VRF1	Enable the exchange of VRF routing information among ISP PE-routers, and enter Address-Family-VRF mode.
(config-router-af)#neighbor 172.3.4.114	
remote-as 300000	Specify the neighbor's (CE1) IP address and ASN value.
(config-router-af)# neighbor 172.3.4.114 activate	Activate neighbor under address family mode
(config-router-af)#neighbor 172.3.4.114 send- community both	Enable extended community attribute for the neighbor.
(config-router-af)#redistribute static	Configure static redistribution.
(config-router-af)#end	Exit Address-Family mode

CE2

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 172.3.4.114/24	Configure the IP address on this interface
(config-if)#exit	Exit interface mode.
(config)#bgp extended-asn-cap	Enable 4-octet ASN capability.
(config)#router bgp 300000	Assign the ASN value (300000) to the router.
(config-router)#neighbor 172.3.4.117 remote-as 200000	Specify the neighbor's IP address (172.3.4.117) and the ASN value of the neighbor (200000).
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 172.3.4.117 activate	Activate neighbor under address family mode
(config-router-af)# exit-address-family	Exit address family mode

Validation

CE1

```
#show ip bgp neighbors
BGP neighbor is 172.4.5.116, remote AS 200000, local AS 100000, external link
```

```
  BGP version 4, local router ID 192.168.52.2, remote router ID 172.4.5.116
  BGP state = Established, up for 00:20:35
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 45 messages, 0 notifications, 0 in queue
  Sent 47 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 172.4.5.115, Local port: 179
Foreign host: 172.4.5.116, Foreign port: 58251
Nexthop: 172.4.5.115
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip bgp vrf all
BGP table version is 3, local router ID is 192.168.52.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   50.1.1.0/24      172.4.5.116           0         100           0       200000 ?
*>   200.1.1.0        172.4.5.116           0         100           0       200000 ?

Total number of prefixes 2


#show ip bgp summary vrf all
BGP router identifier 192.168.52.2, local AS number 100000
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
172.4.5.116            4 200000   46        48      3     0     0  00:21:12
2

Total number of neighbors 1
```

```
Total number of Established sessions 1
```

PE1

```
#show ip bgp neighbors
BGP neighbor is 172.4.5.116, remote AS 200000, local AS 100000, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 172.4.5.116
  BGP state = Established, up for 00:20:35
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 45 messages, 0 notifications, 0 in queue
  Sent 47 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 172.4.5.115, Local port: 179
Foreign host: 172.4.5.116, Foreign port: 58251
Nexthop: 172.4.5.115
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


#show ip bgp vrf all
BGP table version is 3, local router ID is 192.168.52.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   50.1.1.0/24      172.4.5.116           0         100           0       200000 ?
*>   200.1.1.0        172.4.5.116           0         100           0       200000 ?

Total number of prefixes 2

#show ip bgp summary vrf all
BGP router identifier 192.168.52.2, local AS number 100000
BGP table version is 3
1 BGP AS-PATH entries
```

```
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
172.4.5.116              4 200000   46        48      3     0     0  00:21:12
2

Total number of neighbors 1

Total number of Established sessions 1


#clear bgp *
2019 Mar 22 06:16:56.414 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]: Neighbour
[172.4.5.116] Session down due to peer clear
```

PE2

```
#show ip bgp neighbors
BGP neighbor is 172.4.5.116, remote AS 200000, local AS 100000, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 172.4.5.116
  BGP state = Established, up for 00:20:35
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 45 messages, 0 notifications, 0 in queue
  Sent 47 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 172.4.5.115, Local port: 179
Foreign host: 172.4.5.116, Foreign port: 58251
Nexthop: 172.4.5.115
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


#show ip bgp vrf all
BGP table version is 1, local router ID is 172.3.4.117
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path

BGP Route Table for VRF VRF1
*>i  50.1.1.0/24      172.6.7.116          0         100           0        ?
*>   200.1.1.0        0.0.0.0             0         100           32768    ?

Total number of prefixes 2


#show ip bgp summary vrf all
BGP router identifier 192.168.52.2, local AS number 100000
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
172.4.5.116            4 200000   46        48      3     0     0  00:21:12
2

Total number of neighbors 1

Total number of Established sessions 1


#clear bgp *
2019 Mar 22 06:16:56.414 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]: Neighbour
[172.4.5.116] Session down due to peer clear
```

CE2

```
#show ip bgp vrf all
BGP table version is 4, local router ID is 192.168.52.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   50.1.1.0/24      172.3.4.117          0         100           0       200000 ?
*>   200.1.1.0        172.3.4.117          0         100           0       200000 ?

Total number of prefixes 2


#show ip bgp summary vrf all
BGP router identifier 192.168.52.4, local AS number 300000
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
172.3.4.117              4 200000   33        30      4     0     0  00:04:34
2

Total number of neighbors 1

Total number of Established sessions 1


#show ip bgp neighbors
BGP neighbor is 172.3.4.117, remote AS 200000, local AS 300000, external link
  BGP version 4, local router ID 192.168.52.4, remote router ID 172.3.4.117
  BGP state = Established, up for 00:04:40
  Last read 00:00:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 33 messages, 0 notifications, 0 in queue
  Sent 29 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
Local host: 172.3.4.114, Local port: 179
Foreign host: 172.3.4.117, Foreign port: 49361
Nexthop: 172.3.4.114
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:04:40, due to BGP Notification sent
Notification Error Message: (Cease/Administratively Reset.)
```

Nexthop Tracking

Nexthop tracking is used to notify the BGP process asynchronously whenever there is any change in the IGP routes. It reduces the convergence time of BGP routes when IGP routes are changed.

Topology

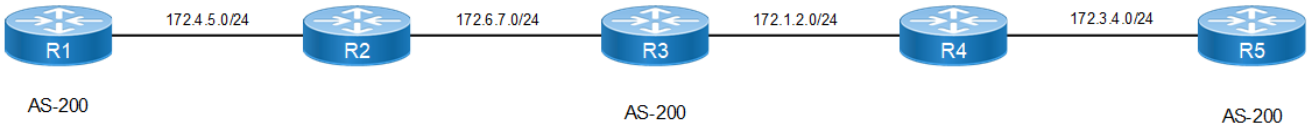

Figure 1-32: BGP Nexthop Tracking

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if)#ip address 100.100.100.100/32 secondary`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the router. The ASN range is <1-65535>.
`(config-router)#neighbor 200.200.200.200 remote-as 200`	Specify the neighbor's IP address (200.200.200.200) and the ASN value of the neighbor (200).
`(config-router)#neighbor 200.200.200.200 update-source lo`	Specify the routing update source.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 200.200.200.200 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#router ospf 1`	Configure the OSPF process (1).
`(config-router)#network 172.4.5.0/24 area 0`	Advertise the network in Area 0.
`(config-router)#redistribute connected`	Redistribute the connected routes.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure the OSPF process (1).
`(config-router)#network 172.4.5.0/24 area 0`	Advertise the network in Area 0.
`(config-router)#network 172.6.7.0/24 area 0`	Advertise the network in Area 0.

R3

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if) #ip address 150.150.150.150/32 secondary`	Configure the IP address on this interface.
`(config-if)#ip address 200.200.200.200/32 secondary`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.

© 2025 IP Infusion Inc. Proprietary

(config)#router bgp 200	Assign the ASN value (200) to the router.
(config-router)#neighbor 100.100.100.100 remote-as 200	Specify the neighbor's IP address (100.100.100.100) and the ASN value of the neighbor (200).
(config-router)#neighbor 100.100.100.100 update-source 200.200.200.200	Specify the routing update source.
(config-router)#neighbor 220.220.220.220 remote-as 200	Specify the neighbor's IP address (220.220.220.220) and the ASN value of the neighbor (200).
(config-router)#neighbor 220.220.220.220 update-source 150.150.150.150	Specify the routing update source.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 100.100.100.100 activate	Activate neighbor under address family mode
(config-router-af)# neighbor 220.220.220.220 activate	Activate neighbor under address family mode
(config-router-af)# exit-address-family	Exit address family mode
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 1	Configure the OSPF process (1).
(config-router)#network 172.6.7.0/24 area 0	Advertise the network in Area 0.
(config-router)#network 172.1.2.0/24 area 0	Advertise the network in Area 0.
(config-router)#redistribute connected	Redistribute the connected routes.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#bgp nexthop-trigger enable	Enable Nexthop tracking.
(config)#bgp nexthop-trigger delay 20	Configure the nexthop trigger-delay time interval.

R4

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure the OSPF process (1).
(config-router)#network 172.1.2.0/24 area 0	Advertise the network in Area 0.
(config-router)#network 172.3.4.0/24 area 0	Advertise the network in Area 0.

R5

#configure terminal	Enter configure mode.
(config)#interface lo	Specify the loopback interface, and enter Interface mode.
(config-if)#ip address 220.220.220.220/32 secondary	Configure the IP address on this interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 200	Assign the ASN value (200) to the router.
(config-router)#neighbor 150.150.150.150 remote-as 200	Specify the neighbor's IP address (150.150.150.150) and the ASN value of the neighbor (200).
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 150.150.150.150 activate	Activate neighbor under address family mode
(config-router-af)# exit-address-family	Exit address family mode

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if)#ip address 220.220.220.220/32 secondary`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config-router)#neighbor 150.150.150.150 update-source lo`	Specify the routing update source.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#router ospf 1`	Configure the OSPF process (1).
`(config-router)#network 172.3.4.0/24 area 0`	Advertise the network in Area 0.
`(config-router)#redistribute connected`	Redistribute the connected routes.

Validation

show ip bgp summary, show ip bgp neighbors, show bgp nexthop-tracking, show ip bgp scan

Nexthop Tracking Delay Timer

The delay interval between routing table walks can be configured for nexthop delay tracking. This time determines how long BGP waits before it starts walking the full BGP routing table after receiving notification from NSM about a next-hop change.

Topology

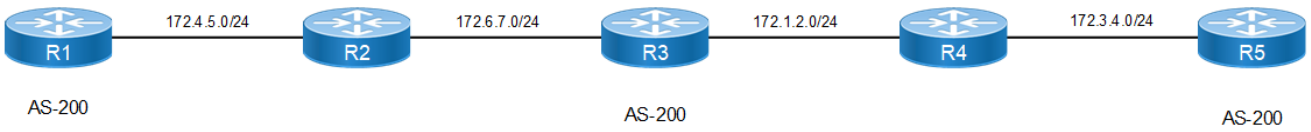

Figure 1-33: Topology for Nexthop Tracking Delay Timer

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if)#ip address 150.150.150.150/32 secondary`	Configure the IP address on this interface.
`(config-if)#ip address 100.100.100.100/32 secondary`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the router. The ASN range is <1-65535>.
`(config-router)#neighbor 200.200.200.200 remote-as 200`	Specify the neighbor's IP address (200.200.200.200) and the ASN value of the neighbor (200).
`(config-router)#neighbor 200.200.200.200 update-source lo`	Specify the routing update source.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode

`(config-router-af)# neighbor 200.200.200.200 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode
`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if)#ip address 150.150.150.150/32 secondary`	Configure the IP address on this interface.
`(config-if)#ip address 100.100.100.100/32 secondary`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the OSPF process (1).
`(config-router)#network 172.4.5.0/24 area 0`	Advertise the network in Area 0.
`(config-router)#redistribute connected`	Redistribute the connected routes.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure the OSPF process (`1`).
`(config-router)#network 172.4.5.0/24 area 0`	Advertise the network in Area `0`.
`(config-router)#network 172.6.7.0/24 area 0`	Advertise the network in Area `0`.

R3

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the loopback interface, and enter Interface mode.
`(config-if)#ip address 200.200.200.200/32`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the router.
`(config-router)#neighbor 100.100.100.100 remote-as 200`	Specify the neighbor's IP address (100.100.100.100) and the ASN value of the neighbor (200).
`(config-router)#neighbor 100.100.100.100 update-source 200.200.200.200`	Specify the routing update source.
`(config-router)#neighbor 220.220.220.220 remote-as 200`	Specify the neighbor's IP address (220.220.220.220) and the ASN value of the neighbor (200).
`(config-router)#neighbor 220.220.220.220 update-source 150.150.150.150`	Specify the routing update source.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# neighbor 100.100.100.100 activate`	Activate neighbor under address family mode
`(config-router-af)# neighbor 220.220.220.220 activate`	Activate neighbor under address family mode
`(config-router-af)# exit-address-family`	Exit address family mode
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#router ospf 1`	Configure the OSPF process (1).
`(config-router)#network 172.6.7.0/24 area 0`	Advertise the network in Area 0.

(config-router)#network 172.1.2.0/24 area 0	Advertise the network in Area 0.
(config-router)#redistribute connected	Redistribute the connected routes.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#bgp nexthop-trigger enable	Enable nexthop tracking.
(config)#bgp nexthop-trigger delay 20	Configure the nexthop trigger-delay time interval.

R4

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure the OSPF process (1).
(config-router)#network 172.1.2.0/24 area 0	Advertise the network in Area 0.
(config-router)#network 172.3.4.0/24 area 0	Advertise the network in Area 0.

R5

#configure terminal	Enter configure mode.
(config)#interface lo	Specify the loopback interface, and enter Interface mode.
(config-if)#ip address 220.220.220.220/32	Configure the IP address on this interface.
(config-if)#exit	Exit interface mode.
(config)#router bgp 200	Assign the ASN value (200) to the router.
(config-router)#neighbor 150.150.150.150 remote-as 200	Specify the neighbor's IP address (150.150.150.150) and the ASN value of the neighbor (200).
(config-router)#neighbor 150.150.150.150 update-source lo	Specify the routing update source.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 150.150.150.150 activate	Activate neighbor under address family mode
(config-router-af)# exit-address-family	Exit address family mode
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 1	Configure the OSPF process (1).
(config-router)#network 172.3.4.0/24 area 0	Advertise the network in Area 0.
(config-router)#redistribute connected	Redistribute the connected routes.

Validation

R1

```
#show ip bgp summary
BGP router identifier 10.12.20.71, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
200.200.200.200          4   200   15         16       1      0      0
00:06:37              0
```

```
Total number of neighbors 1

Total number of Established sessions 1

#show ip bgp neighbors
BGP neighbor is 200.200.200.200, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 10.12.20.71, remote router ID 200.200.200.200
  BGP state = Established, up for 00:06:40
  Last read 00:06:40, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 15 messages, 0 notifications, 0 in queue
  Sent 16 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 100.100.100.100, Local port: 37676
Foreign host: 200.200.200.200, Foreign port: 179
Nexthop: 100.100.100.100
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show bgp nexthop-tracking
Configured NHT: DISABLED
NHT Delay time-interval : 5
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 200, router-id 10.12.20.71

#show ip bgp scan
BGP VRF: (Default) VRF_ID 0
BGP scan interval is 60
scan remain-time: 38
Current BGP nexthop cache:
```

R3

```
#show ip bgp summary
BGP router identifier 200.200.200.200, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
100.100.100.100          4   200   17         19       1      0      0
00:07:41              0
```

```
220.220.220.220          4   200   95        101       1      0      0
00:07:12              0

Total number of neighbors 2

Total number of Established sessions 2

#show ip bgp neighbors
BGP neighbor is 100.100.100.100, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 200.200.200.200, remote router ID 10.12.20.71
  BGP state = Established, up for 00:07:46
  Last read 00:07:46, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 17 messages, 0 notifications, 0 in queue
  Sent 19 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 200.200.200.200
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 200.200.200.200, Local port: 179
Foreign host: 100.100.100.100, Foreign port: 37676
Nexthop: 200.200.200.200
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 220.220.220.220, remote AS 200, local AS 200, internal link
  BGP version 4, remote router ID 220.220.220.220
  local router ID 200.200.200.200
  BGP state = Established, up for 00:07:17
  Last read 00:07:17, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 94 messages, 1 notifications, 0 in queue
  Sent 97 messages, 4 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 150.150.150.150
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 150.150.150.150, Local port: 39831
```

```
Foreign host: 220.220.220.220, Foreign port: 179
Nexthop: 150.150.150.150
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:07:22, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp nexthop-tracking
Configured NHT: ENABLED
NHT Delay time-interval : 20
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 200, router-id 200.200.200.200
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
Total number of IPV4 nexthops : 0
Total number of IPV6 nexthops : 0

#show ip bgp scan
BGP VRF: (Default) VRF_ID 0
BGP scan interval is 60
scan remain-time: 11
Current BGP nexthop cache:
```

R5

```
#show ip bgp summary
BGP router identifier 220.220.220.220, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
150.150.150.150          4   200   99        101       1      0      0
00:08:26              0

Total number of neighbors 1

Total number of Established sessions 1

#show ip bgp neighbors
BGP neighbor is 150.150.150.150, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 220.220.220.220, remote router ID
200.200.200.200
  BGP state = Established, up for 00:08:29
  Last read 00:08:29, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 96 messages, 3 notifications, 0 in queue
  Sent 99 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
 For address family: IPv4 Unicast
```

```
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 220.220.220.220, Local port: 179
Foreign host: 150.150.150.150, Foreign port: 39831
Nexthop: 220.220.220.220
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:08:34, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp nexthop-tracking
Configured NHT: DISABLED
NHT Delay time-interval : 5
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 200, router-id 220.220.220.220

#show ip bgp scan
BGP VRF: (Default) VRF_ID 0
BGP scan interval is 60
scan remain-time: 22
Current BGP nexthop cache:
```

BGP Distance

Administrative distance in BGP can be configured for a specific address family.

Topology

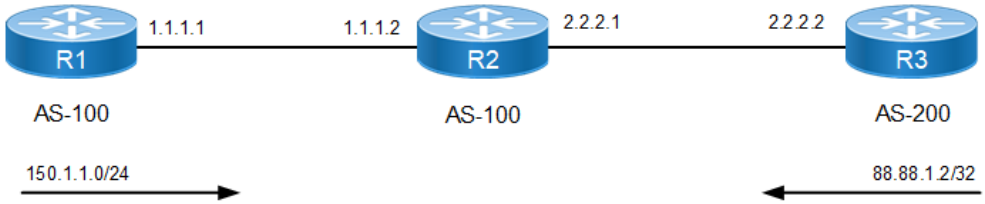

Figure 1-34: Administrative Distance for IPv4 BGP

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter loopback interface mode.
`(config-if)#ip address 150.1.1.1/24 secondary`	Specify IP address for the interface.
`(config-if)#exit`	Exit loopback interface mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the router.

(config-router)#neighbor 1.1.1.2 remote-as 100	Specify the neighbor's IP address and ASN value.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#neighbor 1.1.1.2 activate	Activate the neighbor in address family mode
(config-router-af)#network 150.1.1.0/24	Specify the network to be advertised by the BGP routing process.
(config-router-af)#exit-address-family	Exit address-family ipv4 unicast mode

R2

#configure terminal	Enter configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the router.
(config-router)#neighbor 2.2.2.2 remote-as 200	Specify the neighbor's IP address and ASN value.
(config-router)#neighbor 1.1.1.1 remote-as 100	Specify the neighbor's IP address and the ASN value of another neighbor.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#distance bgp 12 13 120	Configure the administrative distance for external, internal, and local routes received.
(config-router-af)#aggregate-address 150.1.0.0/16 summary-only	Configure a non-AS-set aggregate route on R2. The local distance is applied to this route.
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor in address family mode
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor in address family mode
(config-router-af)#exit-address-family	Exit address-family ipv4 unicast mode

R3

#configure terminal	Enter configure mode.
(config)#interface lo	Enter loopback interface mode.
(config-if)#ip address 88.88.1.2/32 secondary	Specify IP address for the interface.
(config-if)#exit	Exit loopback interface mode.
(config)#router bgp 200	Assign the ASN value (200) to the router.
(config-router)#neighbor 2.2.2.1 remote-as 100	Specify the neighbor's IP address and ASN value.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#network 88.88.1.2/32	Specify the network to be advertised by the BGP routing process.
(config-router-af)#neighbor 2.2.2.1 activate	Activate the neighbor in address family mode
(config-router-af)#exit-address-family	Exitr address-family ipv4 unicast mode

Validation

```
#show ip bgp summary
BGP router identifier 192.168.56.102, local AS number 100
BGP table version is 7
2 BGP AS-PATH entries
```

```
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1.1.1.1                   4   100    8          9       7      0      0
00:02:39               1
2.2.2.2                   4   200    4          4       7      0      0
00:00:38               1

Total number of neighbors 2

Total number of Established sessions 2
#show ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 192.168.52.3, remote router ID 150.1.1.1
  BGP state = Established, up for 00:02:54
  Last read 00:02:54, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 8 messages, 0 notifications, 0 in queue
  Sent 9 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 7, neighbor version 7
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.2, Local port: 49238
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 1.1.1.2
Nexthop global: fe80::a00:27ff:fea6:6e3
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 2.2.2.2, remote AS 200, local AS 100, external link
  BGP version 4, remote router ID 88.88.1.2
  local router ID 192.168.52.3
  BGP state = Established, up for 00:00:53
  Last read 00:00:53, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 4 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 7, neighbor version 7
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes
```

```
 Connections established 1; dropped 0
Local host: 2.2.2.1, Local port: 179
Foreign host: 2.2.2.2, Foreign port: 50072
Nexthop: 2.2.2.1
Nexthop global: fe80::a00:27ff:fe77:264e
Nexthop local: ::
BGP connection: non shared network

#show ip route database bgp
IP Route Table for VRF "default"
B    *> 88.88.1.2/32 [12/0] via 2.2.2.2, eth2, 00:01:19
B    *> 150.1.0.0/16 [120/0] is a summary, Null, 00:02:49
B    *> 150.1.1.0/24 [200/0] via 1.1.1.1, eth1, 00:02:49

Gateway of last resort is not set
#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  1.1.1.0/24 is directly connected, eth1, 00:13:39
C    *>  2.2.2.0/24 is directly connected, eth3, 00:13:04
B    *>  88.88.1.2/32 [12/0] via 2.2.2.2, eth3, 00:06:37
C    *>  127.0.0.0/8 is directly connected, lo, 00:22:15
B    *>  150.1.0.0/16 [120/0] is a summary, Null, 00:11:19
B    *>  150.1.1.0/24 [200/0] via 1.1.1.1, eth1, 00:11:19
C    *>  192.168.52.0/24 is directly connected, eth0, 00:22:13

Gateway of last resort is not set

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     1.1.1.0/24 is directly connected, eth1, 00:17:38
C     2.2.2.0/24 is directly connected, eth3, 00:17:03
B     88.88.1.2/32 [12/0] via 2.2.2.2, eth3, 00:10:36
C     127.0.0.0/8 is directly connected, lo, 00:26:14
B     150.1.0.0/16 [120/0] is a summary, Null, 00:15:18
B     150.1.1.0/24 [200/0] via 1.1.1.1, eth1, 00:15:18
C     192.168.52.0/24 is directly connected, eth0, 00:26:12

Gateway of last resort is not set
```

```
#show ip bgp
BGP table version is 4, local router ID is 192.168.52.3
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf       Weight Path
*>   88.88.1.2/32     2.2.2.2               0         100           0       200
i
*>   150.1.0.0        0.0.0.0               0         100           32768    i
s>i  150.1.1.0/24     1.1.1.1               0         100           0         i

Total number of prefixes 3
```

BGP Weight per Peer

A different weight can be assigned per address family of a peer. For example, a system can be configured to prefer VPN4 routes from peer A and IPv4 routes from peer B.

If the neighbor weight command is given under a specific address-family mode, the peer weight is set for that specific address family. If the address family is not specifically set, the weight is updated for the default address-family.

Topology

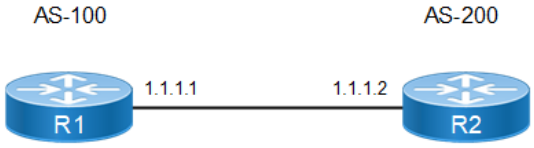

Figure 1-35: BGP Weight Per Peer

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the router.
`(config-router)#neighbor 1.1.1.2 remote-as 200`	Specify the neighbor's IP address and ASN value.
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#neighbor 1.1.1.2 activate`	Activate the neighbor in address family mode
`(config-router-af)#exit-address-family`	Exit address-family ipv4 unicast mode

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the router.
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Specify the neighbor's IP address and ASN value.

(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#neighbor 1.1.1.1 weight 500	Add a weight of 500 to all the routes coming from the neighbor, 1.1.1.1 (only IPv4 routes).
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor in address family mode
(config-router-af)#exit-address-family	Exit address-family ipv4 unicast mode

Validation

R1

```
#show ip bgp summary
BGP router identifier 192.168.56.101, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries
Neighbor V AS MsgRcv MsgSen TblVer InQ OutQ Up/Dow
n State/PfxRcd
1.1.1.2 4 200 6 7 1 0 0 00:02:00
0
Total number of neighbors 1
Total number of Established sessions 1

#show ip bgp neighbors
BGP neighbor is 1.1.1.2, remote AS 200, local AS 100, external link
  BGP version 4, local router ID 192.168.52.2, remote router ID 192.168.52.3
  BGP state = Established, up for 00:01:17
  Last read 00:00:17, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 4 messages, 0 notifications, 0 in queue
  Sent 5 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 179
Foreign host: 1.1.1.2, Foreign port: 34619
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     1.1.1.0/24 is directly connected, eth1, 00:09:10
C     127.0.0.0/8 is directly connected, lo, 00:15:56
C     192.168.52.0/24 is directly connected, eth0, 00:15:52

Gateway of last resort is not set
```

R2

```
#show ip bgp summary
BGP router identifier 192.168.56.102, local AS number 200
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries
Neighbor V AS MsgRcv MsgSen TblVer InQ OutQ Up/Dow
n State/PfxRcd
1.1.1.1 4 100 3 3 1 0 0 00:00:34
0
Total number of neighbors 1
Total number of Established sessions 1


#show ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.2
  BGP state = Established, up for 00:07:14
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 16 messages, 0 notifications, 0 in queue
  Sent 16 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
 Weight500
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.2, Local port: 34619
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 1.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
          i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
          ia - IS-IS inter area, E - EVPN,
          v - vrf leaked
          * - candidate default

    IP Route Table for VRF "default"
    C     1.1.1.0/24 is directly connected, eth1, 00:11:26
    C     127.0.0.0/8 is directly connected, lo, 00:21:36
    C     192.168.52.0/24 is directly connected, eth0, 00:21:32

    Gateway of last resort is not set
```

OSPF as PE-CE Protocol for VPNs

In an VPN environment, customer networks are connected to a VPN-enabled provider backbone. As shown in Figure 1-36, Customer A areas, Areas 0 and 1, are connected to a VPN–enabled provider network. Area 0 and Area 1 have routers CE1 and CE2 running OSPF. MP-iBGP is used between PE1 and PE2 to propagate routes between Site 1 (Area 0) and Site 2 (Area 1). Traditional OSPF-BGP redistribution is performed at PE routers, PE1 and PE2. In this case, routes distributed by CE1 into the MP-iBGP cloud are sent to CE2 as external routes, even though both CE1 and CE2 belong to the same customer.

This behavior can be changed with the additional domain ID configuration. Each VRF should be configured a domain ID on the PE routers. If a PE router gets a route through the MP-iBGP cloud and has to send to any customer site, it checks the domain ID value against the list of stored domain ID values. If the incoming domain ID matches any of the stored IDs, that route is inserted into the customer site with the same type, as it was inserted into the MP-BGP cloud; otherwise, it is inserted as external route.

Topology

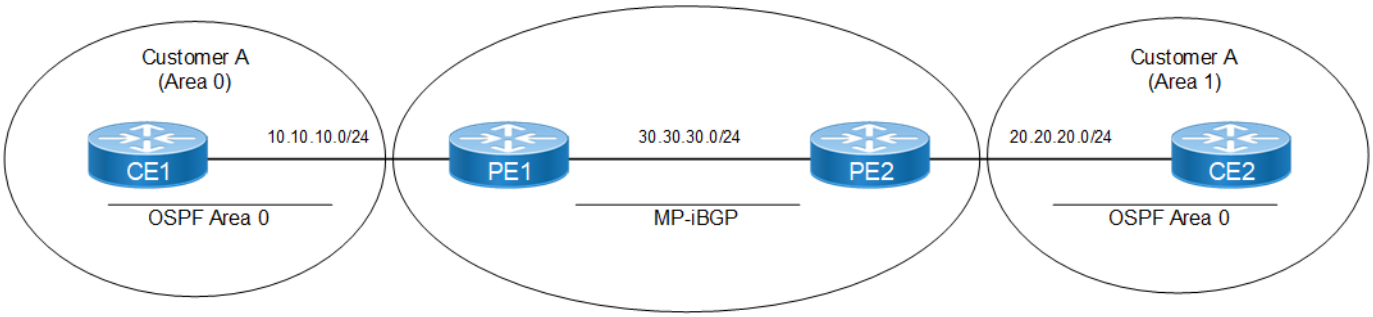

Figure 1-36: OSPF as PE-CE Protocol

Configuration

CE1

`#configure terminal`	Enter configure mode
`(config)#interface lo`	Enter loopback interface mode
`(config-if)#ip address 60.1.1.1/24 secondary`	Specify IP address for the interface
`(config-if)#exit`	Exit loopback interface mode

(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#network 10.10.10.0/24 area 0	Advertise the network in OSPF
(config-router)#network 60.1.1.0/24 area 0	Advertise the loopback IP address in area 0 of router OSPF 1.

PE1

#configure terminal	Enter configure mode.
(config)#ip vrf ABC	Specify the name of the VRF (ABC) to be created.
(config-vrf)#rd 10:100	Assign a route distinguisher (RD) for the VRF, which is a unique value on the router. The RD value can be in ASN:NN or A.B.C.D:NN format.
(config-vrf)#route-target both 10:100	Specify the 2-Octet AS specific or IPv4 specific Transitive Route-Target extended community attribute.
(config-vrf)#exit	Exit VRF mode, and return to Configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip vrf forwarding ABC	Associate interface eth1 to vrf ABC.
(config-if)#exit	Exit interface mode.
(config)#router ospf 1 ABC	Configure OSPF for VRF.
(config-router)#network 10.10.10.0/24 area 0	Advertise the network for OSPF adjacency with CE1.
(config-router)#domain-id 1.1.1.1	Configure the primary domain ID.
(config-router)#domain-id 2.2.2.2 secondary	Configure a secondary domain ID.
(config-router)#domain-id 3.3.3.3 secondary	Configure a secondary domain ID.
(config-router)#exit	Exit Router mode and return to Configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the BGP router.
(config-router)#neighbor 30.30.30.2 remote-as 100	Configure neighbor 30.30.30.2 for iBGP.
(config-router)#address-family vpnv4 unicast	Enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 30.30.30.2 activate	Activate neighbor 30.30.30.2.
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf ABC	Enter Address-Family-VRF mode.
(config-router-af)#redistribute ospf	Specify redistributing routes from OSPF into BGP.
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.

PE2

#configure terminal	Enter configure mode.
(config)#ip vrf ABC	Specify the name of the VRF (ABC) to be created.
(config-vrf)#rd 10:100	Assign a route distinguisher (RD) for the VRF, which is a unique value on the router. The RD value can be in ASN:NN or A.B.C.D:NN format.
(config-vrf)#route-target both 10:100	Specify the 2-Octet AS specific or IPv4 specific Transitive Route-Target extended community attribute.

(config-vrf)#exit	Exit VRF mode, and return to Configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip vrf forwarding ABC	Associate interface eth1 to vrf ABC.
(config-if)#exit	Exit interface mode.
(config)#router ospf 1 ABC	Configure OSPF for VRF.
(config-router)#network 20.20.20.0/24 area 0	Advertise the network for OSPF adjacency with CE1.
(config-router)#domain-id 1.1.1.1	Configure the primary domain ID.
(config-router)#domain-id 2.2.2.2 secondary	Configure a secondary domain ID.
(config-router)#domain-id 3.3.3.3 secondary	Configure a secondary domain ID.
(config-router)#exit	Exit Router mode and return to Configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the BGP router.
(config-router)#neighbor 30.30.30.1 remote-as 100	Configure neighbor 30.30.30.1 for iBGP.
(config-router)#address-family vpnv4 unicast	Enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 30.30.30.1 activate	Activate neighbor 30.30.30.1.
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf ABC	Enter Address-Family-VRF mode.
(config-router-af)#redistribute ospf	Specify redistributing routes from OSPF into BGP.
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.

CE2

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure the routing process, and specify the Process ID (1).
(config-router)#network 20.20.20.0/24 area 0	Advertise the network in OSPF.

Validation

```
#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf      Weight Path
Route Distinguisher: 10:100 (Default for VRF ABC)
*>   10.10.10.0/24    0.0.0.0               2         100           32768    ?
*>   60.1.1.1/32      10.10.10.1            12        100           32768    ?
 Announced routes count = 2
 Accepted routes count = 0
#
```

BGP Multipath for IPv4

BGP supports multipath for IPv4 prefixes. BGP Multipath allows load-balancing traffic among multiple BGP routes. It supports both iBGP and eBGP routes. In case of eBGP, the routes should arrive from same AS number.

Topology

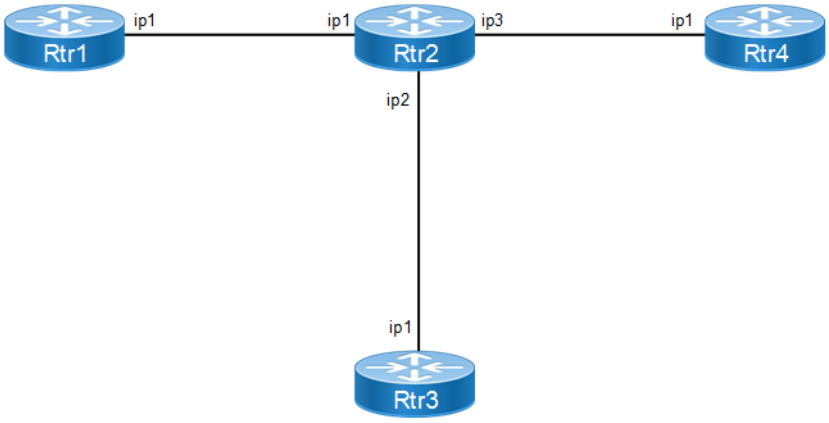

Figure 1-37: Multipath iBGP for IPv4

Configuration

Rtr1

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the BGP router.
`(config-router)#bgp router-id 2.2.2.2`	Configure a fixed Router ID (2.2.2.2).
`(config-router)#neighbor 30.30.30.9 remote-as 100`	Configure neighbor 30.30.30.9 for iBGP.
`(config-router)# address-family ipv4 unicast`	.Under address family, Redistribute the static routes.
`(config-router-af)#redistribute static`	Redistribute the static routes.
`config-router-af)# neighbor 30.30.30.9 activate`	Activate the neighbor
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#exit`	Exit the BGP Router mode and return to the Configure mode.
`(config)#ip route 88.88.0.0/16 Null`	Configure static route.

Rtr3

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the BGP router.
`(config-router)#bgp router-id 4.4.4.4`	Configure a fixed Router ID (4.4.4.4).
`(config-router)#neighbor 40.40.40.9 remote-as 100`	Configure neighbor 40.40.40.9 for iBGP.

(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute static	Redistribute the static routes.
config-router-af)# neighbor 40.40.40.9 activate	Activate the neighbor
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit the BGP Router mode and return to the Configure mode.
(config)#ip route 88.88.0.0/16 Null	Configure static route.

Rtr4

#configure terminal	Enter the Configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the BGP router.
(config-router)#bgp router-id 6.6.6.6	Configure a fixed Router ID (6.6.6.6).
(config-router)#neighbor 50.50.50.9 remote-as 100	Configure neighbor 50.50.50.9 for iBGP.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute static	Redistribute the static routes.
config-router-af)# neighbor 50.50.50.9 activate	Activate the neighbor
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit the BGP Router mode and return to the Configure mode.
(config)#ip route 88.88.0.0/16 Null	Configure static route.

Rtr2

#configure terminal	Enter the Configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the BGP router.
(config)#bgp router-id 9.9.9.9	Configure a fixed Router ID (9.9.9.9).
(config-router)#neighbor 30.30.30.2 remote-as 100	Configure neighbor 30.30.30.2 for iBGP.
(config-router)#neighbor 40.40.40.4 remote-as 100	Configure neighbor 40.40.40.4 for iBGP.
(config-router)#neighbor 50.50.50.6 remote-as 100	Configure neighbor 50.50.50.6 for iBGP.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
config-router-af)# neighbor 30.30.30.2 activate	Activate the neighbor
config-router-af)# neighbor 40.40.40.4 activate	Activate the neighbor
config-router-af)# neighbor 50.50.50.6 activate	Activate the neighbor
config-router-af)# max-paths ibgp 2	Configure iBGP max-paths (2).
(config-router-af)#exit-address-family	Exit address-family mode

Validation

```
#show ip bgp 88.88.0.0
BGP routing table entry for 88.88.0.0/16
Paths: (3 available, best #1, table Default-IP-Routing-Table) Not advertised to any peer
Local
30.30.30.2 from 30.30.30.2 (2.2.2.2)
Origin incomplete, metric 0, localpref 100, valid, internal, multipath- candidate, installed, best
Last update: Wed Mar2 15:17:38 2016

Local
50.50.50.6 from 50.50.50.6 (6.6.6.6)
Origin incomplete, metric 0, localpref 100, valid, internal, multipath- candidate
Last update: Wed Mar2 15:23:58 2016

Local
40.40.40.4 from 40.40.40.4 (4.4.4.4)
Origin incomplete, metric 0, localpref 100, valid, internal, multipath- candidate, installed
Last update: Wed Mar2 15:21:45 2016

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
C   30.30.30.0/24 is directly connected, eth1, 00:15:04 C   40.40.40.0/24 is directly connected, eth6, 00:14:30 C   50.50.50.0/24 is directly connected, eth3, 00:14:46 B 88.88.0.0/16 [200/0] via 40.40.40.4, eth6, 00:02:58
[200/0] via 30.30.30.2, eth1
C  127.0.0.0/8 is directly connected, lo, 00:19:21
C  192.168.52.0/24 is directly connected, eth0, 00:19:16 Gateway of last resort is not set
Gateway of last resort is not set

#show running-config router bgp
!
router bgp 100

bgp router-id 9.9.9.9
max-paths ibgp 2
neighbor 30.30.30.2remote-as100
neighbor 40.40.40.4remote-as100
neighbor 50.50.50.6remote-as100

!
```

```
 address-family ipv4 unicast
 max-paths ibgp 2
  neighbor 30.30.30.2 activate
neighbor 40.40.40.4 activate
neighbor 50.50.50.6 activate
 exit-address-family
!
```

Multipath eBGP

Topology

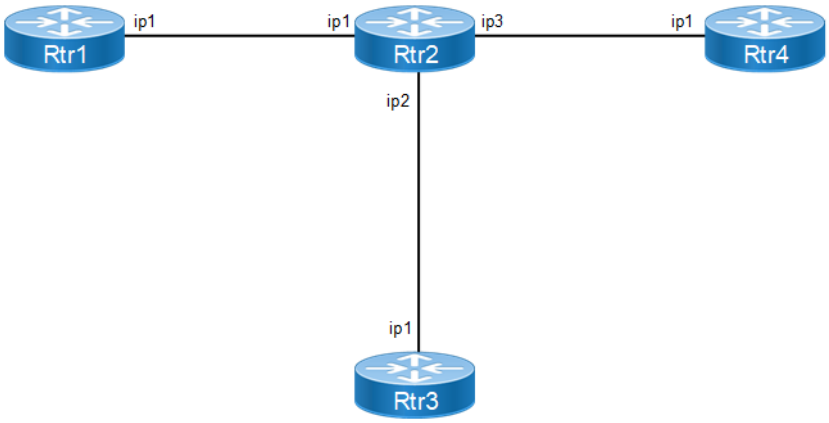

Figure 1-38: Multipath eBGP for IPv4

Configuration

Rtr1

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the BGP router.
`(config-router)#bgp router-id 2.2.2.2`	Configure a fixed Router ID (2.2.2.2).
`(config-router)#neighbor 30.30.30.9 remote-as 100`	Configure neighbor 30.30.30.9 for eBGP.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#redistribute static`	Redistribute the static routes.
`config-router-af)# neighbor 30.30.30.9 activate`	Activate the neighbor
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#exit`	Exit the BGP Router mode and return to the Configure mode.
`(config)#ip route 88.88.0.0/16 Null`	Configure static route.

Rtr3

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the BGP router.
`(config-router)#bgp router-id 4.4.4.4`	Configure a fixed Router ID (4.4.4.4).
`(config-router)#neighbor 40.40.40.9 remote-as 100`	Configure neighbor 40.40.40.9 for eBGP.
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#redistribute static`	Redistribute the static routes.
`config-router-af)# neighbor 40.40.40.9 activate`	Activate the neighbor
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#exit`	Exit the BGP Router mode and return to the Configure mode.
`(config)#ip route 88.88.0.0/16 Null`	Configure static route.

Rtr4

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 200`	Assign the ASN value (200) to the BGP router.
`(config-router)#bgp router-id 6.6.6.6`	Configure a fixed Router ID (6.6.6.6).
`(config-router)#neighbor 50.50.50.9 remote-as 100`	Configure neighbor 50.50.50.9 for eBGP.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#redistribute static`	Redistribute the static routes.
`config-router-af)# neighbor 50.50.50.9 activate`	Activate the neighbor
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config)#exit`	Exit the BGP Router mode and return to the Configure mode.
`(config)#ip route 88.88.0.0/16 Null`	Configure static route.

Rtr2

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the BGP router.
`(config)#bgp router-id 9.9.9.9`	Configure a fixed Router ID (9.9.9.9).
`(config-router)#neighbor 30.30.30.2 remote-as 200`	Configure neighbor 30.30.30.2 for eBGP.
`(config-router)#neighbor 40.40.40.4 remote-as 200`	Configure neighbor 40.40.40.4 for eBGP.
`(config-router)#neighbor 50.50.50.6 remote-as 200`	Configure neighbor 50.50.50.6 for eBGP.
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`config-router-af)# neighbor 30.30.30.2 activate`	Activate the neighbor

© 2025 IP Infusion Inc. Proprietary

config-router-af)# neighbor 40.40.40.4 activate	Activate the neighbor
config-router-af)# neighbor 50.50.50.6 activate	Activate the neighbor
config-router-af)# max-paths ebgp 2	Configure eBGP max-paths (2).
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit the Router mode and return to Configure mode.

Validation

#**show ip bgp 88.88.0.0**

```
BGP routing table entry for 88.88.0.0/16
Paths: (3 available, best #3, table Default-IP-Routing-Table) Advertised to non peer-group peers:
30.30.30.2 50.50.50.6
200
30.30.30.2 from 30.30.30.2 (2.2.2.2)
Origin incomplete metric 0, localpref 100, valid, external, multipath-candidate, installed
Last update: Sat Jan3 02:06:25 1970




200
50.50.50.6 from 50.50.50.6 (6.6.6.6)
Origin incomplete metric 0, localpref 100, valid, external, multipath-candidate Last update: Sat Jan3 02:05:39 1970

200
40.40.40.4 from 40.40.40.4 (4.4.4.4)
Origin incomplete metric 0, localpref 100, valid, external, multipath-candidate, installed, best
Last update: Sat Jan3 02:05:11 1970

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
C   30.30.30.0/24 is directly connected, eth1, 05:26:26
C   40.40.40.0/24 is directly connected, eth6, 05:25:52
C   50.50.50.0/24 is directly connected, eth3, 05:26:08
B   88.88.0.0/16 [20/0] via 40.40.40.4, eth6, 00:01:38
    [20/0] via 30.30.30.2, eth1
C   127.0.0.0/8 is directly connected, lo, 05:30:43
```

```
C   192.168.52.0/24 is directly connected, eth0, 05:30:38
Gateway of last resort is not set #show running-config router bgp
!
router bgp 100
bgp router-id 9.9.9.9
max-paths ebgp 2
neighbor 30.30.30.2 remote-as 200
neighbor 40.40.40.4 remote-as 200
neighbor 50.50.50.6 remote-as 200
!
 address-family ipv4 unicast
 neighbor 30.30.30.2 activate
 neighbor 40.40.40.4 activate
 neighbor 50.50.50.6 activate
 exit-address-family
!
```

Multipath eiBGP

Use this command to set the number of equal-cost multi-path (ECMP) routes for both eBGP and iBGP. This feature allows to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol.

When enabled, this feature can perform load balancing on eBGP and/or iBGP paths. With multiple eBGP and iBGP paths, selection is per all paths sorted according to BGP rules and hence either of all eBGP or iBGP or both will be marked as candidates for multipath programming i.e., it is possible either all eBGP paths or all iBGP paths or both eBGP and iBGP are programmed as multipaths.

Exceptions during path selections when eiBGP is enabled:

- AS Path length check is ignored.
- Origin Check is ignored.
- Peer type check is ignored.

IGP Metric check is ignored.

Topology

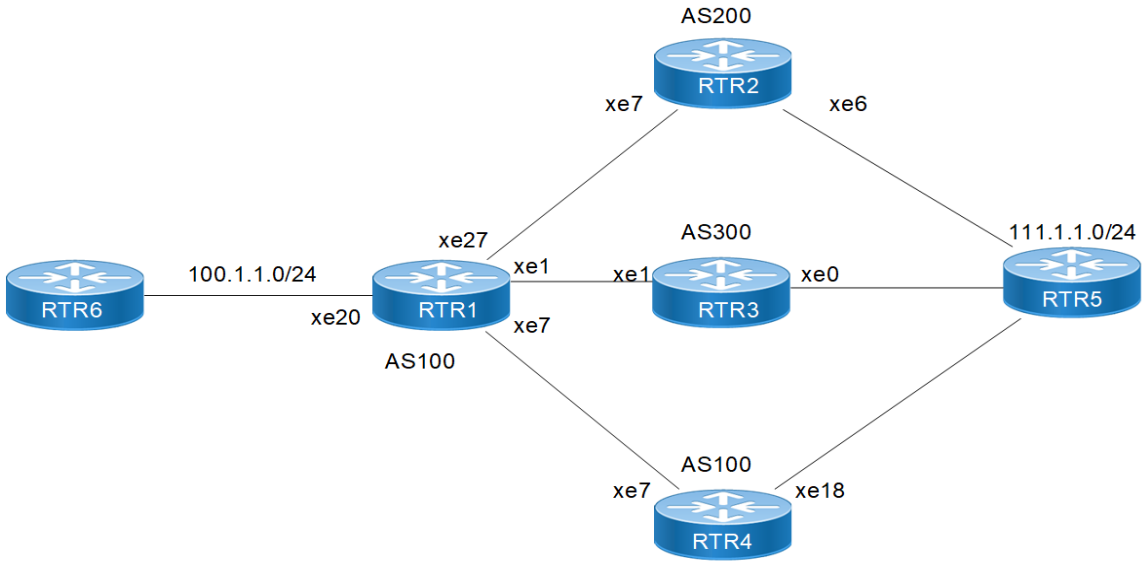

Figure 1-39: Multipath eiBGP topology

Configuration

RTR1

`#configure terminal`	Enter Configure mode.
`(config)#interface lo`	Enter Interface mode
`(config-if)# ip address 41.41.41.41/32 secondary`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe27`	Enter Interface mode
`(config-if)# ip address 21.1.1.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe1`	Enter Interface mode
`(config-if)# ip address 22.1.1.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe7`	Enter Interface mode
`(config-if)# ip address 23.1.1.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe20`	Enter Interface mode
`(config-if)# ip address 100.1.1.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)# router bgp 100`	Enter Router BGP mode
`(config-router)# neighbor 21.1.1.2 remote-as 200`	Define BGP neighbors. 21.1.1.2 is the IP address of the neighbor (RTR2) and 200 is the neighbors AS number

(config-router)# neighbor 22.1.1.2 remote-as 300	Define BGP neighbors. 22.1.1.2 is the IP address of the neighbor (RTR3) and 300 is the neighbors AS number
(config-router)# neighbor 23.1.1.2 remote-as 100	Define BGP neighbors. 23.1.1.2 is the IP address of the neighbor (RTR4) and 100 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#max-paths eibgp 4	Configure eiBGP max-paths (4).
(config-router-af)#redistribute connected	Redistribute connected routes into BGP
(config-router-af)#neighbor 21.1.1.2 activate	Activate the neighbor
(config-router-af)#neighbor 22.1.1.2 activate	Activate the neighbor
(config-router-af)#neighbor 23.1.1.2 activate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

RTR2

#configure terminal	Enter Configure mode.
(config)#interface lo	Enter Interface mode
(config-if)# ip address 45.45.45.45/32 secondary	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe7	Enter Interface mode
(config-if)# ip address 21.1.1.2/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe6	Enter Interface mode
(config-if)# ip address 111.1.1.11/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)# router bgp 200	Enter Router BGP mode
(config-router)# neighbor 21.1.1.1 remote-as 100	Define BGP neighbors. 21.1.1.1 is the IP address of the neighbor (RTR1) and 100 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#redistribute connected	Redistribute connected routes
(config-router-af)#neighbor 21.1.1.1 activate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

RTR3

#configure terminal	Enter Configure mode.
(config)#interface lo	Enter Interface mode
(config-if)# ip address 42.42.42.42/32 secondary	Assign IP address to interface

(config-if)#exit	Exit interface mode
(config)#interface xe0	Enter Interface mode
(config-if)# ip address 111.1.1.42/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe1	Enter Interface mode
(config-if)# ip address 22.1.1.2/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)# router bgp 300	Enter Router BGP mode
(config-router)# neighbor 22.1.1.1 remote-as 100	Define BGP neighbors. 22.1.1.1 is the IP address of the neighbor (RTR1) and 100 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#redistribute connected	Redistribute connected routes
(config-router-af)#neighbor 22.1.1.1 activate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

RTR4

#configure terminal	Enter Configure mode.
(config)#interface lo	Enter Interface mode
(config-if)# ip address 43.43.43.43/32 secondary	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe7	Enter Interface mode
(config-if)# ip address 23.1.1.2/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe18	Enter Interface mode
(config-if)# ip address 111.1.1.43/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)# router bgp 100	Enter Router BGP mode
(config-router)# neighbor 23.1.1.1 remote-as 100	Define BGP neighbors. 23.1.1.1 is the IP address of the neighbor (RTR1) and 100 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#redistribute connected	Redistribute connected routes
(config-router-af)#neighbor 23.1.1.1 activate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

Validation

RTR1

```
# sh ip bgp summary
BGP router identifier 15.1.1.2, local AS number 100
BGP table version is 4
3 BGP AS-PATH entries
0 BGP community entries
4  Configured ebgp ECMP multipath: Currently set at 4

Neighbor    V   AS   MsgRcv  MsgSen TblVer   InQ   OutQ    Up/Down State/PfxRcd
21.1.1.2    4   200   13       16       4      0      0  00:04:28             3
22.1.1.2    4   300   10       14       4      0      0  00:03:05             4
23.1.1.2    4   100    7        9       4      0      0  00:01:49             4

Total number of neighbors 3

Total number of Established sessions 3

#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            11.1.1.0/24 [200/0] via 23.1.1.2, xe7, 00:11:26
B            11.11.11.11/32 [20/0] via 21.1.1.2, xe27, 00:10:14
B            17.1.1.0/24 [20/0] via 22.1.1.2, xe1, 00:10:54
C            21.1.1.0/24 is directly connected, xe27, 00:47:36
C            22.1.1.0/24 is directly connected, xe1, 00:49:33
C            23.1.1.0/24 is directly connected, xe7, 00:49:11
C            41.41.41.41/32 is directly connected, lo, 01:30:34
B            42.42.42.42/32 [20/0] via 22.1.1.2, xe1, 00:10:54
B            43.43.43.43/32 [200/0] via 23.1.1.2, xe7, 00:11:26
B            45.45.45.45/32 [20/0] via 21.1.1.2, xe27, 00:10:56
C            100.1.1.0/24 is directly connected, xe20, 00:08:31
B            111.1.1.0/24 [200/0] via 23.1.1.2, xe7, 00:11:26
                          [200/0] via 22.1.1.2, xe1,
                          [200/0] Via 21.1.1.2, xe27
C            127.0.0.0/8 is directly connected, lo, 01:39:19

Gateway of last resort is not set

#sh interface counters rate mbps
```

```
--------------------------------------------------------------------------

Interface

Rx mbps

Rx pps

Tx mbps

Tx pps

--------------------------------------------------------------------------
xe1                  0.00            0               264.30          22024
xe7                  0.00            0               254.86          21238
xe20                 822.35          68529           0.02            3
xe27                 0.00            0               215.04          17920
```

6BGP AS-PATH Multipath-relax

BGP will not load balance across multiple paths by default. We can configure it to do so with the `max-paths ebgp <no-of-multipaths>` command. The criterion of this command is that all attributes must match (Weight, Local preference, AS Path, etc). This is acceptable if we are multi-homed to a single AS, but what if we are multi-homed to different AS.

BGP AS PATH multipath relax effectively allows for ECMP to be done across different neighboring ASN's.

Topology

Below topology explains about BGP AS PATH multipath relax functionality.

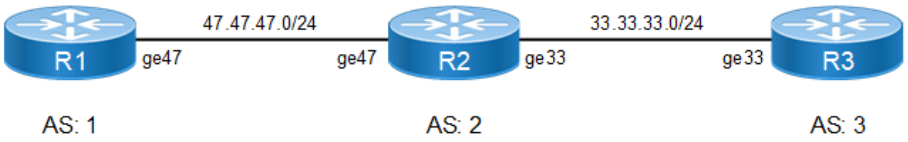

Figure 1-40: BGP AS-PATH Multipath-relax Topology

Configuration

R1

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter Interface loopback
`(config-if)#ip address 100.1.1.1/24 secondary`	Configure IP address for interface
`(config-if)#exit`	Exit interface mode
`(config)#interface ge47`	Enter Interface loopback

(config-if)#ip address 47.47.47.1/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter the Configure mode
(config)#router bgp 1	Assign the ASN value (1) to the BGP router
(config-router)#neighbor 47.47.47.2 remote-as 2	Configure eBGP neighbor.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#neighbor 47.47.47.2 activate	Activate neighbor
(config-router-af)#network 100.1.1.0/24	Advertise the loopback network into BGP.
(config-router-af)#end	Exit from router BGP and address-family config mode

R2

#configure terminal	Enter the Configure mode.
(config)#interface ge33	Enter interface mode
(config-if)#ip address 33.33.33.2/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter the Configure mode
(config)#interface ge47	Enter interface mode
(config-if)#ip address 47.47.47.2/24	Configure IP address for interface.
(config-if)#exit	Exit from interface mode and enter the Configure mode
(config)#router bgp 2	Assign the ASN value (2) to the BGP router.
(config-router)#neighbor 33.33.33.3 remote-as 3	Configure eBGP neighbor.
(config-router)#neighbor 47.47.47.1 remote-as 1	Configure eBGP neighbor.
(config-router)#address-family ipv4 unicast	Enter the address family IPv4 unicast mode
(config-router-af)# neighbor 33.33.33.3 activate	Activate the neighbor
config-router-af)# neighbor 47.47.47.1 activate	Activate the neighbor
(config-router-af)# max-paths ebgp 8	Configure eBGP Multipath.
(config-router-af)#exit	Exit from address family mode and enter the Router configure mode
(config-router)# bgp bestpath as-path multipath-relax	Configure BGP AS PATH Multipath relax.
(config-router-af)#end	End the address family mode.

R3

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter Interface loopback.
(config-if)#ip address 100.1.1.1/24 secondary	Configure IP address for interface.
(config-if)#exit	Exit from interface mode and enter the Configure mode.
(config)#interface ge33	Enter Interface loopback.

(config-if)#ip address 33.33.33.3/24	Configure IP address for interface.
(config-if)#exit	Exit from interface mode and enter the Configure mode.
(config)#router bgp 3	Assign the ASN value (3) to the BGP router.
(config-router)#neighbor 33.33.33.2 remote-as 2	Configure eBGP neighbor.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# neighbor 33.33.33.2 activate	Activate the neighbor
(config-router-af)#network 100.1.1.0/24	Advertise the loopback network into BGP.
(config-router-af)#end	Exit from router BGP and address-family config mode

Validation

```
R2#show running-config bgp
!
router bgp 2
bgp bestpath as-path multipath-relax max-paths ebgp 8
neighbor 33.33.33.3 remote-as 3
neighbor 47.47.47.1 remote-as 1
!
 address-family ipv4 unicast
 neighbor 33.33.33.3 activate
 neighbor 47.47.47.1  activate
 max-paths ebgp 8
 bgp bestpath as-path multipath-relax
 exit-address-family
!
R2#show ip bgp 100.1.1.0
BGP routing table entry for 100.1.1.0/24
Paths: (2 available, best #1, table Default-IP-Routing-Table) Advertised to non peer-group peers:
47.47.47.1
3
33.33.33.3 from 33.33.33.3 (33.33.33.3)
Origin IGP, metric 0, localpref 100, valid, external, multipath- candidate, installed, best
Last update: Tue Feb 23 03:13:14 2016
1
47.47.47.1 from 47.47.47.1 (62.57.1.1)
Origin IGP, metric 0, localpref 100, valid, external, multipath-
candidate, installed
Last update: Tue Feb 23 03:13:15 2016

R2#show ip bgp summary
BGP router identifier 192.168.52.3, local AS number 2 BGP table version is 2
2 BGP AS-PATH entries


0 BGP community entries
```

```
8   Configured ebgp ECMP multipath: Currently set at 8
1   Configured ibgp ECMP multipath: Currently set at 1
1   Configured eibgp ECMP multipath: Currently set at 1

NeighborVASMsgRcvMsgSen TblVerInQOutQUp/ DownState/PfxRcd
33.33.33.3
00:01:101
47.47.47.1
00:06:331

Total number of neighbors 2

Total number of Established sessions 2

R2#show ip bgp
BGP table version is 2, local router ID is 192.168.52.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

NetworkNext HopMetricLocPrfWeightPath
*> 100.1.1.0/2447.47.47.1010001 i
*   33.33.33.301000  3 i

Total number of prefixes 1
```

BGP FIB Install (Selective Route Download)

The BGP—Selective Route Download feature allows a network administrator to selectively download some or none of the BGP routes into the Routing Information Base (RIB). The primary application for this feature is to suppress the unnecessary downloading of certain BGP routes to the RIB or Forwarding Information Base (FIB) on a dedicated route reflector, which propagates BGP updates without carrying transit traffic. The feature thereby helps to maximize resources available and to improve routing scalability and convergence on the dedicated route reflector.

With RFC 4456, the concept of route reflection was defined; this would allow configuring designated one or more BGP routers in iBGP network as route reflectors. BGP relaxes the re-advertising restriction on these route reflectors, allowing them to accept and propagate IBGP routes to their clients.

The role of a dedicated route reflector (RR) is to propagate BGP updates without participating in the actual forwarding of transit traffic. That means the RR does not need to have all BGP routes downloaded into its RIB or FIB. It is beneficial for the RR to preserve its resources by not processing and storing those routes.

By default, BGP routes are downloaded to the RIB. To save resources on a dedicated route reflector, such downloading can be reduced or prevented by configuring a table map. A table map is so named because it controls what is put into the BGP routing table.

By reducing the route installation in the dedicated route reflectors, we can maximize availability of resources and improve routing scalability and convergence.

A new command ‘table map’ is being introduced to achieve this. A table map controls what is put into the BGP routing table. When configured it would reduce or prevent downloading routes to RIB.

Table map command references ‘route map’ rules available in BGP to control the routes going into the BGP routing table.

Table-map command can be used in two ways:

- When a simple table-map command is given (without filter option), the route map referenced in the table-map command shall be used to set certain properties (such as the traffic index) of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.
- When the option ‘filter’ is given in the table map command, the route map referenced is used to control whether a BGP route is to be downloaded to the IP RIB (hence the filter). A BGP route is not downloaded to the RIB if it is denied by the route map.

Topology

Below topology explains about BGP FIB Install functionality

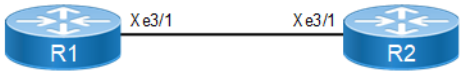

Figure 1-41: BFP FIB Install Topology

Configuration

R1

`#configure terminal`	Enter the Configure mode.
`(config)#interface xe3/1`	Enter interface mode.
`(config-if)#ip address 20.1.1.1/24`	Configure IP address for interface
`(config-if)#exit`	Exit interface mode
`(config)# router bgp 100`	Assign the ASN value (100) to the BGP router.
`(config-router)#neighbor 20.1.1.2 remote-as 100`	Configure neighbor in IBGP
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#redistribute static`	Redistribute static routes to advertise to its neighbor
`(config-router-af)#neighbor 20.1.1.2 activate`	Activate the neighbor
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#exit`	Exit Router mode and enter Configure mode
`(config)#ip route 1.1.1.0/24 xe3/1`	Configure static route.
`(config)#ip route 2.2.2.0/24 xe3/1`	Configure static route.
`(config)#ip route 3.3.3.0/24 xe3/1`	Configure static route.
`(config)#ip route 4.4.4.0/24 xe3/1`	Configure static route.
`(config)#ip route 5.5.5.0/24 xe3/1`	Configure static route.
`(config)#ip route 6.6.6.0/24 xe3/1`	Configure static route.

R2

`#configure terminal`	Enter the Configure mode.
`(config)#interface xe3/1`	Configure IP address for interface

(config-if)#ip address 20.1.1.2/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter into Configure mode
(config)#router bgp 100	Assign the ASN value (100) to the BGP router.
(config-router)#neighbor 20.1.1.1 remote-as 100	Configure neighbor iBGP.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute static	Redistribute the static routes.
(config-router-af)#neighbor 20.1.1.1 activate	Activate the neighbor
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit the BGP Router mode and return to the Configure mode.
(config)#ip access-list 1	Login to Configure access-list parameters
(config-ip-acl)#permit ipip 2.2.2.0 0.0.0.225 any	Configure access-list by allowing only one route to install in FIB table.
(config-ip-acl)#exit	Exit assess list mode
(config)# route-map test permit 1	Configure route-map to match access-list
(config-route-map)# match ip address 1	Match the above configured access-list 1
(config-route-map)#exit	Exit from route-map Configure mode and enter into Configure mode
(config)#router bgp 100	Enter into BGP router mode
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# table-map test filter	Apply table-map with route-map created and with filter option
(config-router-af)#end	Exit from router and Configure mode

Validation

Table-map with Filter Option

Verify BGP neighborship is up between R1 and R2. Before applying table-map in R2, all routes will be installed in FIB table, as in below output.

R1

```
#show ip bgp summary
BGP router identifier 20.1.1.1, local AS BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries
number 100
NeighborVASMsgRcvMsgSen TblVerInQOutQUp/Dow
n   State/PfxRcd
20.1.1.241005
6   3
0
0
00:01:31
0
Total number of neighbors 1
Total number of Established sessions 1
```

```
#show ip bgp
BGP table version is 1, local router ID is 192.168.52.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l -
labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete


   NetworkNext HopMetricLocPrfWeight Path
*> 1.1.1.0/240.0.0.0010032768?
*> 2.2.2.0/240.0.0.0010032768?
*> 3.3.3.0/240.0.0.0010032768?
*> 4.4.4.0/240.0.0.0010032768?
*> 5.5.5.0/240.0.0.0010032768?
*> 6.6.6.0/240.0.0.0010032768?

Total number of prefixes 6 #

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
S   1.1.1.0/24[1/0]isdirectlyconnected,eth1,00:06:54
S   2.2.2.0/24[1/0]isdirectlyconnected,eth1,00:06:35
S   3.3.3.0/24[1/0]isdirectlyconnected,eth1,00:06:26
S   4.4.4.0/24[1/0]isdirectlyconnected,eth1,00:06:17
S   5.5.5.0/24[1/0]isdirectlyconnected,eth1,00:06:09
S   6.6.6.0/24[1/0]isdirectlyconnected,eth1,00:06:01
C   20.1.1.0/24 is directly connected, eth1, 00:07:32 C127.0.0.0/8 is directly
connected, lo, 00:08:21
C   192.168.52.0/24 is directly connected, eth0, 00:08:17

Gateway of last resort is not set #
```

R2

```
#show ip bgp
BGP table version is 1, local router ID is 192.168.52.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l -
labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   NetworkNext HopMetricLocPrfWeight Path
*> 1.1.1.0/240.0.0.0010032768?
*> 2.2.2.0/240.0.0.0010032768?
*> 3.3.3.0/240.0.0.0010032768?
*> 4.4.4.0/240.0.0.0010032768?
```

```
*> 5.5.5.0/240.0.0.0010032768?
*> 6.6.6.0/240.0.0.0010032768?


Total number of prefixes 6 #

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
B   1.1.1.0/24[200/0]via20.1.1.1,eth1,00:13:44
B   2.2.2.0/24[200/0]via20.1.1.1,eth1,00:13:44
B   3.3.3.0/24[200/0]via20.1.1.1,eth1,00:13:44
B   4.4.4.0/24[200/0]via20.1.1.1,eth1,00:13:44
B   5.5.5.0/24[200/0]via20.1.1.1,eth1,00:13:44
B   6.6.6.0/24[200/0]via20.1.1.1,eth1,00:13:44
C   20.1.1.0/24 is directly connected, eth1, 00:14:12 C127.0.0.0/8 is directly
connected, lo, 00:25:26
C   192.168.52.0/24 is directly connected, eth0, 00:25:23

Gateway of last resort is not set #
```

Table-map With Filter Option

```
Now verify after applying table-map with filter option, only one route will be installed
in FIB table according to route-map and access-list configured, BGP table remains same,
table-map effect will be seen only for FIB table.
After applying table-map, clear BGP with "clear ip bgp table-map."
(config)#router bgp 100
(config-router)#address-family ipv4-unicast (config-router-af)#table-map test filter
(config-router-af)#end
#clear ip bgp table-map

#show ip bgp
BGP table version is 2, local router ID is 192.168.52.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l -
labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   NetworkNext HopMetricLocPrfWeightPath
*>i 1.1.1.0/2420.1.1.101000?
*>i 2.2.2.0/2420.1.1.101000?
*>i 3.3.3.0/2420.1.1.101000?
*>i 4.4.4.0/2420.1.1.101000?
```

```
*>i 5.5.5.0/2420.1.1.101000?
*>i 6.6.6.0/2420.1.1.101000?

Total number of prefixes 6 #

#show ip bgp summary
BGP router identifier 192.168.52.5, local AS number 100 BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries

NeighborVASMsgRcvMsgSen TblVerInQOutQUp/DownState/ PfxRcd
20.1.1.141004039 2   0   0   00:18:33
6

Total number of neighbors 1

Total number of Established sessions 1 #

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
B   2.2.2.0/24 [200/0] via 20.1.1.1, eth1, 00:00:26
C   20.1.1.0/24 is directly connected, eth1, 00:19:01 C127.0.0.0/8 is directly
connected, lo, 00:30:15
C   192.168.52.0/24 is directly connected, eth0, 00:30:12

Gateway of last resort is not set
```

Table-map Without Filter Option

Remove filter option while applying table-map as below in R2

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Enter into BGP router mode
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# table-map test`	Apply table-map with route-map created and with filter option
`(config-router-af)#end`	Exit from router and configure mode

```
#show running-config bgp
!
router bgp 100 redistribute static
```

```
neighbor 20.1.1.1 remote-as 100 table-map test
!
 address-family ipv4 unicast
 neighbor 20.1.1.1 activate
 exit-address-family
!
#clear ip bgp table-map

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
B   1.1.1.0/24[200/0]via20.1.1.1,eth1,00:00:04
B   2.2.2.0/24[200/0]via20.1.1.1,eth1,00:00:04
B   3.3.3.0/24[200/0]via20.1.1.1,eth1,00:00:04
B   4.4.4.0/24[200/0]via20.1.1.1,eth1,00:00:04
B   5.5.5.0/24[200/0]via20.1.1.1,eth1,00:00:04
B   6.6.6.0/24[200/0]via20.1.1.1,eth1,00:00:04
C   20.1.1.0/24 is directly connected, eth1, 00:31:16 C127.0.0.0/8 is directly
connected, lo, 00:42:30
C   192.168.52.0/24 is directly connected, eth0, 00:42:27

Gateway of last resort is not set #
```

Note: Same can be tried with IPV6 VRF–v4 and VRF–v6 address-families.

BGP Best Path Selection Process

BGP assigns the first valid path as the current best path. BGP then compares the best path with the next path in the list, until BGP reaches the end of the list of valid paths. Below steps provides the rules that are used to determine the best path:

1. Prefer the path with the highest WEIGHT.
2. Prefer the path with the highest LOCAL_PREF.
3. Prefer the path that was locally originated via a network or aggregate BGP subcommand or through redistribution from an IGP.
4. Prefer the path with the shortest AS_PATH.

Note: Beware of these items:

- This step is skipped if user has configured the `bgp bestpath as-path ignore` command.
- If `bgp bestpath compare-confed-aspath` is configured then Prefer the path with the shortest AS_CONFED path.

5. Prefer the path with the lowest ORIGIN type.

Note: Beware of below listed items:

- IGP is lower than Exterior Gateway Protocol (EGP), and EGP is lower than INCOMPLETE.

6. Prefer the path with the lowest multi-exit discriminator (MED).

Note: Beware of these items:

- By default, MED is compared in these cases:
  - MEDs are compared only if the first AS in the AS_SEQUENCE is the same for multiple paths
  - If both the paths are internal as routes
  - If paths have confederation as-path then MEDs are compared only if the first AS in the BGP_AS_CONFED_SEQUENCE is the same for multiple paths
  - To override all above checks, user can configure `bgp always-compare-med` command

7. Prefer eBGP over iBGP paths.

Note: Beware of below listed item:

- EBGP is preferred over IGBP or EBGP is preferred over CONFED.

8. Path learned from LU Address-family is preferred over IPv4 Unicast Address-family.

Note: Beware of these items:

- This is Exception Rule for IPv4 Labeled-Unicast Address-family.
- This rule applicable only for IPv4 Labeled-Unicast/Unicast routes over default VRF.

9. Prefer the path with the lowest IGP metric to the BGP next hop.

10. Determine if multiple paths require installation in the routing table for BGP Multipath and mark the ECMP candidate.

11. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path won't displace an older one, even if it was the preferred route based on the additional decision criteria below. This has to be enabled by BGP command `bestpath tie-break-on-age`

Note: Beware of these items:

- Skip this step if any of these items are true:
  - If `bgp bestpath compare-routerid` is configured in addition to `bestpath tie-break-on-age`, then this step will be skipped.
  - If the router ID is same for multiple paths, because the routes were received from the same router, then this step will be skipped.

12. Router ID and Originator Id:

- If `bgp bestpath compare-routerid` is configured, then prefer the route that comes from the BGP router with the lowest Router ID.
- If `bgp bestpath dont-compare-originator-id` is not configured, prefer the route that comes from the BGP router with the lowest Router ID.

Note: Beware of the below listed item:

- If a path contains Route Reflector (RR) attributes, the Originator ID is substituted for the Router ID in the path selection process.

- If `bgp bestpath dont-compare-originator-id` is configured, prefer the route that comes from the BGP router with the lowest router ID. In this case, Originator ID is not compared even if the RR attribute is present.

13. If the originator or Router ID is the same for multiple paths, prefer the path with the minimum cluster list length.

Prefer the path that comes from the lowest neighbor address.

BGP Dampening

BGP supports route dampening for IPv4 and IPv6 prefixes. Route dampening minimizes the instability caused by route flapping. A penalty is added for every flap in a flapping route. As soon as the total penalty reaches the suppress limit, the advertisement of the route is suppressed. This penalty is decayed according to the configured half time value. Once the penalty is lower than the reuse limit, the route advertisement is unsuppressed. The dampening information is purged from the router once the penalty becomes less than half of the reuse limit.

Topology

In this example, a successful TCP connection is being established between the routers.

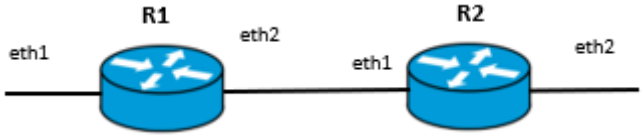

Figure 1-42: BGP dampening

IPv4 Configuration

R1

`#configure terminal`	Enter configure mode
`(config)#interface lo`	Enter loopback interface mode
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the secondary loopback address
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 101.1.0.1/24`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 100`	Configure BGP with the AS number 100
`(config-router)#neighbor 10.1.1.2 remote-as 200`	Define the BGP neighbor, and establish a TCP session. 10.1.1.2 is the IP address of one of the neighbors (R2), and 200 is the neighbor's AS number.

(config-router)#neighbor 100.1.0.2 remote-as 300	Define the BGP neighbor, and establish a TCP session. 100.1.0.2 is the IP address of one of the neighbors on interface eth1, and 300 is the neighbor's AS number.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute connected	Enable redistribute connected
(config-router-af)#neighbor 10.1.1.2 activate	Activate the neighbor
(config-router-af)#neighbor 100.1.0.2 activate	Activate the neighbor
(config-router-af)#exit-address-family	Exit address-family mode

R2

#configure terminal	Enter configure mode
(config)#interface lo	Enter loopback interface mode
(config-if)# ip address 2.2.2.2/32 secondary	Configure the secondary loopback address
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 10.1.1.2/24	Configure the IP address of the interface
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode
(config-if)#ip address 101.1.0.1/24	Configure the IP address of the interface
(config-if)#exit	Exit interface mode
(config)#router bgp 200	Configure BGP with the AS number 100
(config-router)#neighbor 10.1.1.1 remote-as 100	Define the BGP neighbor, and establish a TCP session. 10.1.1.1 is the IP address of one of the neighbors (R1), and 100 is the neighbor's AS number.
(config-router)#neighbor 101.1.0.2 remote-as 400	Define the BGP neighbor, and establish a TCP session.101.1.0.2 is the IP address of one of the neighbors on eth2 interface, and 400 is the neighbor's AS number.
(config-router)#neighbor 100.1.0.2 remote-as 300	Define the BGP neighbor, and establish a TCP session. 100.1.0.2 is the IP address of one of the neighbors of router R1 on eth1 interface, and 300 is the neighbor's AS number.
(config-router)#neighbor 100.1.0.2 ebgp-multihop 2	Increase BGP neighbors with ebgp-multihop value
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#neighbor 10.1.1.1 activate	Activate BGP neighbor
(config-router-af)#neighbor 101.1.0.2 activate	Activate BGP neighbor

(config-router-af)#redistribute connected	Enable redistribute connected
(config-router-af)#bgp dampening	Enable BGP dampening with default values: • Reachability half-life is 15 minutes • Reuse limit is 750: • Suppress limit is 2000 • Max-suppress value is 60 minutes • Un-reachability half-life is 15 minutes

Validation

R2

Verify the BGP dampening parameters.

```
#show ip bgp dampening parameters

dampening 15 750 2000 60 15 Dampening Control Block(s):
Reachability Half-Life time : 15 min Reuse penalty: 750
Suppress penalty: 2000
Max suppress time: 60 min Un-reachability Half-Life time : 15 min Max penalty (ceil):
11999
Min penalty (floor): 375
```

Verify BGP dampened paths for flapping networks.

```
#show ip bgp dampening dampened-paths
BGP table version is 21, local router ID is 4.4.4.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l -
labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   NetworkFromReusePath
d  200.1.0.010.1.1.100:29:00 100300i
d  200.2.0.010.1.1.100:28:20 100300i
d  200.3.0.010.1.1.100:28:20 100300i
d  200.4.0.010.1.1.100:28:20 100300i
d  200.5.0.010.1.1.100:28:20 100300i
d  200.6.0.010.1.1.100:28:20 100300i
d  200.7.0.010.1.1.100:28:20 100300i
d  200.8.0.010.1.1.100:28:20 100300i
d  200.9.0.010.1.1.100:28:20 100300i
d  200.10.0.010.1.1.100:28:20 100300i
```

Verify BGP dampening flap statistics for flapping networks.

```
#show ip bgp dampening flap-statistics
BGP table version is 21, local router ID is 4.4.4.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l -
labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   NetworkFromFlapsDurationReusePath
d  200.1.0.010.1.1.1800:18:3700:29:10100 300i
d  200.2.0.010.1.1.1700:14:2200:29:00100 300i
d  200.3.0.010.1.1.1700:14:2200:29:00100 300i
d  200.4.0.010.1.1.1700:14:2200:29:00100 300i
d  200.5.0.010.1.1.1700:14:2200:29:00100 300i
d  200.6.0.010.1.1.1700:14:2200:29:00100 300i
d  200.7.0.010.1.1.1700:14:2200:29:00100 300i
d  200.8.0.010.1.1.1700:14:2200:29:00100 300i
d  200.9.0.010.1.1.1700:14:2200:29:00100 300i
d  200.10.0.010.1.1.1700:14:2200:29:00100 300i
```

IPv6 Configuration

R1

`#configure terminal`	Enter configure mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 address 2000:0:0:1::1/64`	Configure the IPv6 address of the interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 2000:0:2:1::1/64`	Configure the IPv6 address of the interface
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 100`	Configure BGP with the AS number 100
`(config-router)#neighbor 2000:0:0:1::2 remote-as 300`	Define the BGP neighbor, and establish a TCP session. 2000:0:0:1::2 is the IP address of one of the neighbors on interface eth1, and 300 is the neighbor's AS number.
`(config-router)#neighbor 2000:0:2:1::2 remote-as 200`	Define the BGP neighbor, and establish a TCP session. 2000:0:2:1::2 is the IP address of one of the neighbors (R2), and 200 is the neighbor's AS number.
`(config-router)#address-family ipv6 unicast`	Enter IPv6 address family
`(config-router)#redistribute connected`	Enable redistribute connected
`(config-router-af)neighbor 2000:0:0:1::2 activate`	Activate BGP neighbor
`(config-router-af)neighbor 2000:0:2:1::2 activate`	Activate BGP neighbor

R2

`#configure terminal`	Enter configure mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 2000:0:2:1::2/64`	Configure the IPv6 address of the interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 2000:0:1:1::1/64`	Configure the IPv6 address of the interface

`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Configure BGP with the AS number 200
`(config-router)#address-family ipv6 unicast`	Enter IPv6 address family
`(config-router-af)#redistribute connected`	Enable redistribute connected
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#neighbor 2000:0:1:1::2 remote-as 400`	Define the BGP neighbor, and establish a TCP session. 2000:0:1:1::2 is the IP address of one of the neighbors on interface eth2, and 400 is the neighbor's AS number.
`(config-router)#neighbor 2000:0:2:1::1 remote-as 100`	Define the BGP neighbor, and establish a TCP session. 2000:0:2:1::1 is the IP address of one of the neighbors (R1), and 100 is the neighbor's AS number.
`(config-router)#address-family ipv6 unicast`	Enter IPv6 address-family
`(config-router-af)#bgp dampening`	Enable BGP dampening with default values: • Reachability half-life is 15 minutes • Reuse limit is 750 • Suppress limit is 2000 • Max-suppress value is 60 minutes • Un-reachability half-life is 15 minutes
`(config-router-af)#neighbor 2000:0:1:1::2 activate`	Activate BGP neighbor
`(config-router-af)#neighbor 2000:0:2:1::1 activate`	Activate BGP neighbor

Validation

R2

Verify the IPv6 BGP dampening parameters.

```
#sh bgp dampening parameters

 dampening 15 750 2000 60 15
 Dampening Control Block(s):
  Reachability Half-Life time    : 15 min
  Reuse penalty                  : 750
  Suppress penalty               : 2000
  Max suppress time              : 60 min
  Un-reachability Half-Life time : 15 min
  Max penalty (ceil)             : 11999
  Min penalty (floor)            : 375
```

Verify IPv6 BGP dampened paths for flapping networks.

```
#sh bgp dampening dampened-paths
BGP table version is 7, local router ID is 4.4.4.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network            From                  Reuse    Path
*d   3000:0:1:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:2:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:3:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:4:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:5:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:6:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:7:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:8:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:9:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
*d   3000:0:a:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                                    00:18:30 100 300 i
```

Verify IPv6 BGP dampening flap statistics for flapping networks.

```
#sh bgp dampening flap-statistics
BGP table version is 7, local router ID is 4.4.4.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network            From                Flaps  Duration  Reuse    Path
*d   3000:0:1:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:2:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:3:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:4:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:5:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:6:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:7:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:8:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:9:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
*d   3000:0:a:1::/64  2000:0:2:1::1(fe80::ba6a:97ff:fed6:23d4)
                                            4  00:05:19  00:18:30  100 300 i
```

BGP Authentication

BGP authentication allows users to receive selected routing information, enhancing security of their network traffic. When BGP authentication is enabled on a router, the router verifies routing packets it receives by exchanging a password that is configured on both the sending and receiving routers.

In this example, both R1 and R2 have ABC as the password. Configure the same password on all routers that are to communicate using BGP in a network.

Topology

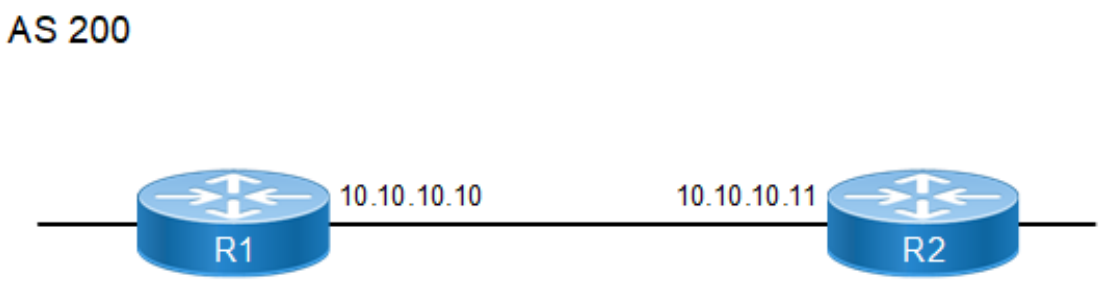

Figure 1-43: BGP Authentication

Configuration

R1

`#configure terminal`	Enter configure mode
`(config)#interface xe0`	Enter interface mode
`(config-if)#ip address 10.10.10.10/24`	Assign ip address
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Enter BGP router mode
`(config-router)#neighbor 10.10.10.11 remote-as 300`	Configure neighborship
`(config-router)#neighbor 10.10.10.11 authentication-key 0 ABC`	Configure authentication for BGP neighbors
`(config-router)#address-family ipv4 unicast`	Enter ipv4 address family
`(config-router-af)#neighbor 10.10.10.11 activate`	Activate neighborship
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#end`	End config mode

© 2025 IP Infusion Inc. Proprietary

R2

`#configure terminal`	Enter configure mode
`(config)#interface ce2/1`	Enter interface mode
`(config-if)#ip address 10.10.10.11/24`	Assign ip address
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Enter BGP router mode
`(config-router)#neighbor 10.10.10.10 remote-as 200`	Configure neighborship
`(config-router)#neighbor 10.10.10.10 authentication-key 0 ABC`	Configure authentication for BGP neighbors
`(config-router)#address-family ipv4 unicast`	Enter ipv4 address family
`(config-router-af)#neighbor 10.10.10.10 activate`	Activate neighborship
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#end`	End config mode

Validation

R1

```
#show running-config bgp
!
router bgp 200
 neighbor 10.10.10.11 remote-as 300
 neighbor 10.10.10.11 authentication-key 0x624ac41428f81e33
 !
 address-family ipv4 unicast
 neighbor 10.10.10.11 activate
 exit-address-family
!
#show ip bgp neighbors
BGP neighbor is 10.10.10.11, remote AS 300, local AS 200, external link
  BGP version 4, local router ID 22.22.22.22, remote router ID 2.2.2.2
  BGP state = Established, up for 00:09:14
  Last read 00:00:19, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 23 messages, 0 notifications, 0 in queue
  Sent 24 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.10, Local port: 179
```

```
Foreign host: 10.10.10.11, Foreign port: 37590
Nexthop: 10.10.10.10
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

R2

```
#show running-config bgp
!
router bgp 300
 neighbor 10.10.10.10 remote-as 200
 neighbor 10.10.10.10 authentication-key 0x624ac41428f81e33
 !
 address-family ipv4 unicast
 neighbor 10.10.10.10 activate
 exit-address-family
!
#show ip bgp neighbors
BGP neighbor is 10.10.10.10, remote AS 200, local AS 300, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 22.22.22.22
  BGP state = Established, up for 00:13:57
  Last read 00:00:22, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 33 messages, 0 notifications, 0 in queue
  Sent 34 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.11, Local port: 37590
Foreign host: 10.10.10.10, Foreign port: 179
Nexthop: 10.10.10.11
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

BGP Unnumbered

This section contains configurations for BGP unnumbered interface which provides BGP peering with minimal configuration.

Overview

BGP protocol is used to exchange IP prefixes between AS. For BGP neighbor ship to be established, IPv4 address configuration on peer is pre-requisite. In a large network, this can consume a lot of your address space, requiring a separate IP address for each peer-facing interface apart from administrator effort in configuration. When a BGP peer advertises an IPv4 prefix, it must include an IPv4 next hop address, which is usually the address of the advertising router; for this each BGP peer should have an IPv4 address. This feature is to enable BGP peering with minimal configuration, less IPv4 address-space.

For DC use-case, where hundreds of switches can be connected in CLOS topology, configuring each neighbor is both time consuming and (IPv4) address hungry.

To minimize this, BGP unnumbered can, avoid the need for an IP address on each BGP interface and by removing the need to configure the IP address and ASN of each neighbor. This feature uses link local ipv6 address of interface as per RFC-5549.

Note: Do not configure additional link-local addresses on the unnumbered interface as the BGP Unnumbered session comes up on the default Link-local Address.

Note: If BGP unnumbered is configured with only if /30 or /31 on the IPv4 address interface, then the BGP session is established through the IPv4 address. To establish the BGP session through link-local, configure `no ipv6 nd supress-ra` and `ipv6 nd ra-interval <value>` commands on the interface without any IP address.

Note: To enable router solicitation or advertisement on an interface for un-numbered support, the `no ipv6 nd suppress-ra` command must be configured on both sides of the interface. Currently, suppression is disabled by default. This facilitates learning a neighbor's IPv6 Link-Local Address (LLA), which is necessary for sending OPEN messages to establish BGP neighborship.

Note: As per RFC 5549, extended next-hop is not supported when IPv6 prefixes are advertised to a next-hop device running IPv4. Both devices must have either IPv6 or IPv4 addresses. AFI 1, SAFI 1, and IPv6 next-hop (Link-Local Address) are the only supported combinations with this feature. Global Unicast Address (GUA) for IPv6 next-hop is not supported.

Note: The hardware supports only data path packet forwarding (for IPv4 prefixes received through BGP neighborship with the unnumbered interface over IPv6 LLA). Ping and traceroute do not work since entries are not installed in the kernel.

Topology

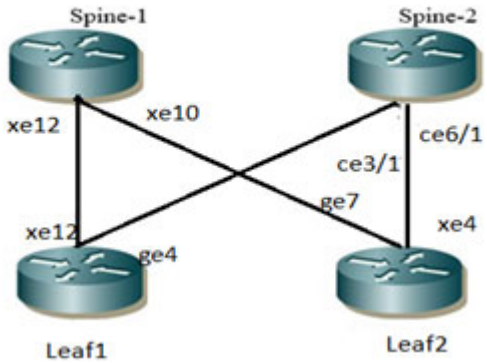

Figure 1-44: BGP-Unnumbered

Configuration

Spine 1

`configure terminal`	Enter configure mode
`(config)#interface lo`	Enter interface mode for loopback interface
`(config-if)#ip add 1.1.1.1/32 secondary`	Assign secondary interface to loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface xe12`	Enter interface mode
`(config-if)# ipv6 nd ra-interval 4`	Assign the IPv6 Router Advertisements interval
`(config-if)#exit`	Exit interface mode
`(config)#interface xe10`	Enter interface mode
`(config-if)# ip address 10.10.10.2/31`	Assign IP address to the interface in /31 subnet
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 100`	Enter bgp router mode
`(config-router)#bgp router-id 1.1.1.1`	Assign router id for BGP
`(config-router)#bgp unnumbered-mode`	Enter bgp unnumbered mode
`(config-router-unnum)#neighbor xe12 remote-as internal`	Configure iBGP neighborship
`(config-router-unnum)#neighbor xe10 remote-as internal`	Configure iBGP neighborship
`(config-router-unnum)# exit-unnumbered-mode`	Exit unnumbered mode
`(config-router)#address-family ipv4 unicast`	Enter address family mode for IPv4 unicast
`(config-router-af)# bgp v4-unnumbered-mode`	Enter unnumbered mode under ipv4 unicast af
`(config-router-v4-unnum)# neighbor xe10 activate`	Activate the neighbor
`(config-router-v4-unnum)# neighbor xe12 activate`	Activate the neighbor
`(config-router-v4-unnum)# exit-v4-unnumbered-mode`	Exit unnumbered mode under ipv4 unicast af
`(config-router-af)# exit-address-family`	Exit address family mode for IPv4 unicast
`(config-router)#end`	End Config mode

Spine 2

`configure terminal`	Enter configure mode
`(config)#interface lo`	Enter interface mode for loopback interface
`(config-if)# ip address 2.2.2.2/32 secondary`	Assign secondary interface to loopback
`(config-if)#exit`	Exit interface mode
`(config)# interface ce6/1`	Enter interface mode
`(config-if)# ipv6 nd ra-interval 4`	Assign the IPv6 Router Advertisements interval
`(config-if)#exit`	Exit interface mode
`(config)# interface ce3/1`	Enter interface mode

© 2025 IP Infusion Inc. Proprietary

(config-if)# ip address 20.20.20.2/31	Assign IP address to the interface in /31 subnet
(config-if)#exit	Exit interface mode
(config)#router bgp 100	Enter bgp router mode
(config-router)#bgp router-id 2.2.2.2	Assign router id for BGP
(config-router)#bgp unnumbered-mode	Enter bgp unnumbered mode
(config-router-unnum)#neighbor ce3/1 remote-as internal	Configure iBGP neighborship
(config-router-unnum)#neighbor ce6/1 remote-as internal	Configure iBGP neighborship
(config-router-unnum)# exit-unnumbered-mode	Exit unnumbered mode
(config-router)#address-family ipv4 unicast	Enter address family mode for IPv4 unicast
(config-router-af)# bgp v4-unnumbered-mode	Enter unnumbered mode under ipv4 unicast af
(config-router-v4-unnum)# neighbor ce3/1 activate	Activate the neighbor
(config-router-v4-unnum)# neighbor ce6/1 activate	Activate the neighbor
(config-router-v4-unnum)# exit-v4-unnumbered-mode	Exit unnumbered mode under ipv4 unicast af
(config-router-af)# exit-address-family	Exit address family mode for IPv4 unicast
(config-router)#end	End Config mode

Leaf 1

configure terminal	Enter configure mode
(config)#interface lo	Enter interface mode for loopback interface
(config-if)# ip address 3.3.3.3/32 secondary	Assign secondary interface to loopback
(config-if)#exit	Exit interface mode
(config)#interface xe12	Enter interface mode
(config-if)# ipv6 nd ra-interval 4	Assign the IPv6 Router Advertisements interval
(config-if)#exit	Exit interface mode
(config)#interface ge4	Enter interface mode
(config-if)# ip address 10.10.10.3/31	Assign IP address to the interface in /31 subnet
(config-if)#exit	Exit interface mode
(config)#router bgp 100	Enter bgp router mode
(config-router)#bgp router-id 3.3.3.3	Assign router id for BGP
(config-router)#bgp unnumbered-mode	Enter bgp unnumbered mode
(config-router-unnum)#neighbor xe12 remote-as internal	Configure iBGP neighborship
(config-router-unnum)#neighbor ge4 remote-as internal	Configure iBGP neighborship
(config-router-unnum)# exit-unnumbered-mode	Exit unnumbered mode
(config-router)#address-family ipv4 unicast	Enter address family mode for IPv4 unicast
(config-router-af)# bgp v4-unnumbered-mode	Enter unnumbered mode under ipv4 unicast af

(config-router-v4-unnum)# neighbor ge4 activate	Activate the neighbor
(config-router-v4-unnum)# neighbor xe12 activate	Activate the neighbor
(config-router-v4-unnum)# exit-v4-unnumbered-mode	Exit unnumbered mode under ipv4 unicast af
(config-router-af)# exit-address-family	Exit address family mode for IPv4 unicast
(config-router)#end	End Config mode

Leaf 2

configure terminal	Enter configure mode
(config)#interface lo	Enter interface mode for loopback interface
(config-if)#ip address 4.4.4.4/32 secondary	Assign secondary interface to loopback
(config-if)#exit	Exit interface mode
(config)#interface xe4	Enter interface mode
(config-if)# ipv6 nd ra-interval 4	Assign the IPv6 Router Advertisements interval
(config-if)#exit	Exit interface mode
(config)#interface ge7	Enter interface mode
(config-if)# ip address 10.10.10.2/31	Assign IP address to the interface in /31 subnet
(config-if)#exit	Exit interface mode
(config)#router bgp 100	Enter bgp router mode
(config-router)#bgp router-id 4.4.4.4	Assign router id for BGP
(config-router)#bgp unnumbered-mode	Enter bgp unnumbered mode
(config-router-unnum)#neighbor xe4 remote-as internal	Configure iBGP neighborship
(config-router-unnum)#neighbor ge7 remote-as internal	Configure iBGP neighborship
(config-router-unnum)# exit-unnumbered-mode	Exit unnumbered mode
(config-router)#address-family ipv4 unicast	Enter address family mode for IPv4 unicast
(config-router-af)# bgp v4-unnumbered-mode	Enter unnumbered mode under ipv4 unicast af
(config-router-v4-unnum)# neighbor ge7 activate	Activate the neighbor
(config-router-v4-unnum)# neighbor xe4 activate	Activate the neighbor
(config-router-v4-unnum)# exit-v4-unnumbered-mode	Exit unnumbered mode under ipv4 unicast af
(config-router-af)# exit-address-family	Exit address family mode for IPv4 unicast
(config-router)#end	End Config mode

Validation

Spine 1:

```
Spine1#show ip bgp neighbors
BGP neighbor is 10.10.10.3, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 1.1.1.1, remote router ID 4.4.4.4
  BGP state = Established, up for 00:22:12
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 54 messages, 0 notifications, 0 in queue
  Sent 54 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.2, Local port: 179
Foreign host: 10.10.10.3, Foreign port: 49242
Nexthop: 10.10.10.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is fe80::eac5:7aff:fefe:97e0, remote AS 100, local AS 100, internal
link
  BGP version 4, local router ID 1.1.1.1, remote router ID 3.3.3.3
  BGP state = Established, up for 00:29:54
  Last read 00:00:12, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
```

```
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 71 messages, 0 notifications, 0 in queue
  Sent 74 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: fe80::eac5:7aff:fe8b:a82a, Local port: 179
Foreign host: fe80::eac5:7aff:fefe:97e0, Foreign port: 37116
Nexthop: 1.1.1.1
Nexthop global: fe80::eac5:7aff:fe8b:a82a
Nexthop local: fe80::eac5:7aff:fe8b:a82a
BGP connection: shared network
```

Spine 2:

```
Spine2#show ip bgp neighbors
BGP neighbor is 20.20.20.3, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 00:21:15
  Last read 00:00:12, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 52 messages, 0 notifications, 0 in queue
  Sent 51 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
```

```
 Connections established 1; dropped 0
Local host: 20.20.20.2, Local port: 59380
Foreign host: 20.20.20.3, Foreign port: 179
Nexthop: 20.20.20.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is fe80::36ef:b6ff:fe31:dd3f, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 4.4.4.4
  BGP state = Established, up for 00:29:31
  Last read 00:00:21, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 70 messages, 0 notifications, 0 in queue
  Sent 72 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: fe80::ce37:abff:fe3f:9f63, Local port: 179
Foreign host: fe80::36ef:b6ff:fe31:dd3f, Foreign port: 33368
Nexthop: 2.2.2.2
Nexthop global: fe80::ce37:abff:fe3f:9f63
Nexthop local: fe80::ce37:abff:fe3f:9f63
BGP connection: shared network
```

Leaf 1

```
Leaf1#show ip bgp neighbors
BGP neighbor is 20.20.20.2, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 2.2.2.2
  BGP state = Established, up for 00:21:32
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 52 messages, 0 notifications, 0 in queue
  Sent 54 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
```

```
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 20.20.20.3, Local port: 179
Foreign host: 20.20.20.2, Foreign port: 59380
Nexthop: 20.20.20.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is fe80::eac5:7aff:fe8b:a82a, remote AS 100, local AS 100, internal
 link
  BGP version 4, local router ID 3.3.3.3, remote router ID 1.1.1.1
  BGP state = Established, up for 00:30:46
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 74 messages, 0 notifications, 0 in queue
  Sent 73 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: fe80::eac5:7aff:fefe:97e0, Local port: 37116
Foreign host: fe80::eac5:7aff:fe8b:a82a, Foreign port: 179
Nexthop: 3.3.3.3
Nexthop global: fe80::eac5:7aff:fefe:97e0
Nexthop local: fe80::eac5:7aff:fefe:97e0
BGP connection: shared network
```

Leaf 2

```
Leaf2#show ip bgp neighbors
BGP neighbor is 10.10.10.2, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 4.4.4.4, remote router ID 1.1.1.1
  BGP state = Established, up for 00:23:24
  Last read 00:00:09, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 56 messages, 0 notifications, 0 in queue
  Sent 57 messages, 0 notifications, 0 in queue
```

```
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 10.10.10.3, Local port: 49242
Foreign host: 10.10.10.2, Foreign port: 179
Nexthop: 10.10.10.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is fe80::ce37:abff:fe3f:9f63, remote AS 100, local AS 100, internal
 link
  BGP version 4, local router ID 4.4.4.4, remote router ID 2.2.2.2
  BGP state = Established, up for 00:30:09
  Last read 00:00:08, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 72 messages, 0 notifications, 0 in queue
  Sent 72 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: fe80::36ef:b6ff:fe31:dd3f, Local port: 33368
Foreign host: fe80::ce37:abff:fe3f:9f63, Foreign port: 179
Nexthop: 4.4.4.4
Nexthop global: fe80::36ef:b6ff:fe31:dd3f
Nexthop local: fe80::36ef:b6ff:fe31:dd3f
BGP connection: shared network
```

BGP Blackhole Community Attribute

A blackhole route is used to forward unwanted or undesirable traffic into a black hole. In other words, a special logical interface called a null interface, is used to create the black hole. Static routes are created for destinations that are not

desirable, and the static route configuration points to the null interface. Any traffic that has a destination address that has a best match of the black hole static route automatically will be dropped.

Note: 65535:666 is reserved for Blackhole community.

Topology

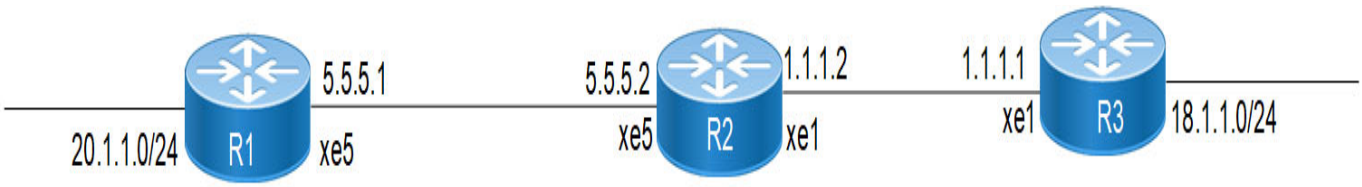

Figure 1-45: BGP Blackhole Community Attribute topology

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)#interface xe5`	Enter Interface mode
`(config-if)# ip address 5.5.5.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe20`	Enter Interface mode
`(config-if)# ip address 20.1.1.1/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)# router bgp 100`	Enter Router BGP mode
`(config-router)# neighbor 5.5.5.2 remote-as 200`	Define BGP neighbors. 5.5.5.2 is the IP address of the neighbor (R2) and 200 is the neighbors AS number
`(config-router)# address-family ipv4 unicast`	Enter into BGP address family IPv4
`(config-router-af)#neighbor 5.5.5.2 activate`	Activate the neighbor
`(config-router-af)#network 20.1.1.0/24`	Advertise networks with prefix
`(config-router-af)# commit`	Commit the configurations
`(config-router-af)# end`	Return to privilege mode

R2

`#configure terminal`	Enter Configure mode.
`(config)#interface xe5`	Enter Interface mode
`(config-if)# ip address 5.5.5.2/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)#interface xe1`	Enter Interface mode
`(config-if)# ip address 1.1.1.2/24`	Assign IP address to interface
`(config-if)#exit`	Exit interface mode
`(config)# router bgp 200`	Enter Router BGP mode

(config-router)# neighbor 5.5.5.1 remote-as 100	Define BGP neighbors. 5.5.5.1 is the IP address of the neighbor (R1) and 100 is the neighbors AS number
(config-router)# neighbor 1.1.1.1 remote-as 300	Define BGP neighbors. 1.1.1.1 is the IP address of the neighbor (R3) and 100 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#neighbor 5.5.5.1 activate	Activate the neighbor
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

R3

#configure terminal	Enter Configure mode.
(config)#interface xe1	Enter Interface mode
(config-if)# ip address 1.1.1.1/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)#interface xe18	Enter Interface mode
(config-if)# ip address 18.1.1.1/24	Assign IP address to interface
(config-if)#exit	Exit interface mode
(config)# router bgp 300	Enter Router BGP mode
(config-router)# neighbor 1.1.1.2 remote-as 200	Define BGP neighbors. 1.1.1.2 is the IP address of the neighbor (R2) and 200 is the neighbors AS number
(config-router)# address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#neighbor 1.1.1.2 activate	Activate the neighbor
(config-router-af)#network 18.1.1.0/24	Advertise networks with prefix
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

Black Hole configuration on R3

#configure terminal	Enter Configure mode.
(config)#route-map D permit 10	Enter Route-map mode to set the match operation
(config-route-map)#set community no-export 65535:666 additive	Configure Reserved Black hole community in Route-map mode
(config-route-map)#commit	Commit the configuration
(config-route-map)#exit	Return to configuration mode
(config)#router bgp 300	Enter Router BGP mode
(config-router)#address-family ipv4 unicast	Enter into BGP address family IPv4
(config-router-af)#neighbor 1.1.1.2 route-map D out	Apply Route-map for the neighbor 1.1.1.2 in out direction
(config-router-af)#commit	Commit the configurations
(config-router-af)#end	Return to privilege mode
#clear ip bgp * soft out	Soft reset after applying Route-map

Validation

R2

```
# show ip bgp community
BGP table version is 4, local router ID is 5.5.5.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best,
i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf   Weight Path
*>   18.1.1.0/24      1.1.1.1               0        100       0   300 i

Total number of prefixes 1

#show ip bgp 18.1.1.0/24
BGP routing table entry for 18.1.1.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table, not advertised
to EBGP peer)
  Not advertised to any peer
  AS path:300
   Nexthop:1.1.1.1 from 1.1.1.1 (Remote Id:1.1.1.1)
      Origin IGP, metric 0, localpref 100      valid, external, best, source
safi: 1
      Community: 65535:666 no-export
      Not advertised to any peer
      Last update: Tue Apr 16 21:48:01 2019

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, xe1, 00:10:22
C            5.5.5.0/24 is directly connected, xe5, 00:10:49
B            18.1.1.0/24 [20/0] is a summary, Null, 00:02:00
B            20.1.1.0/24 [20/0] via 5.5.5.1, xe5, 00:05:46
C            127.0.0.0/8 is directly connected, lo, 00:35:31

Gateway of last resort is not set
```

R1

```
#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
```

```
       * - candidate default

IP Route Table for VRF "default"
C            5.5.5.0/24 is directly connected, xe5, 00:15:41
C            20.1.1.0/24 is directly connected, xe20, 00:14:06
C            127.0.0.0/8 is directly connected, lo, 00:37:28

Gateway of last resort is not set
```

CHAPTER 2 BGP Graceful Restart Configuration

During a BGP restart, all BGP peers detect that a session had gone down and come back up. OcNOS invalidates the associated portion of the IP forwarding cache, does a BGP route re-computation, and generates BGP routing updates. The forwarding tables become corrupted and unstable.

Graceful restart helps minimize these negative effects on routing caused by a BGP restart by allowing the restarting BGP router to temporarily retain routing information and continue forwarding packets while BGP restarts. In this way, even while a router rebuilds routing and forwarding tables, the router continues to operate across the TCP connection.

Graceful restart allows a restarting router, and its neighbors, to continue forwarding packets, without disrupting network performance. Because neighboring routers assist in the restart, the restarting router can quickly resume full operation.

The graceful restart capability extends to the case when a configuration change forces a peer reset.

Graceful reset is a refinement of graceful restart to help ensure smooth restarts when a configuration change forces BGP peer reset.

Topology

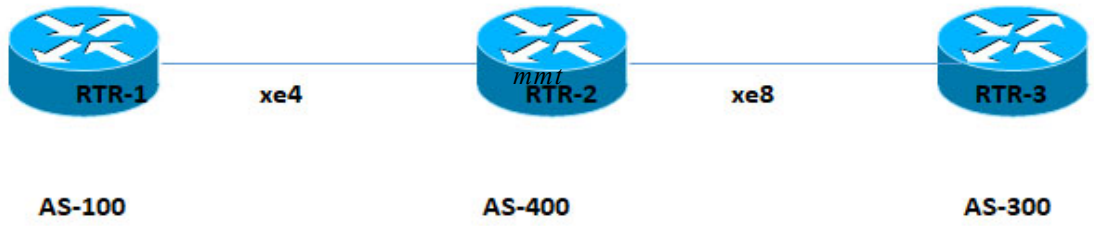

Figure 2: Device topology for BGP in VR/VRF

RTR1

`#configure terminal`	Enter Configuration mode.
`(config)#interface lo`	Enter interface mode for loopback.
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure ip address on loopback.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe4`	Enter interface mode for eth2.
`(config-if)#ip address 10.10.10.1/24`	Configure ip address on eth2.
`(config-if)#exit`	Exit interface mode for eth2.
`(config)# router bgp 100`	Enter router bgp mode.
`(config-router)# bgp router-id 1.1.1.11.1.1.1`	Configure bgp router-id same as loopback ip address.
`(config-router)# bgp graceful-restart`	Configure Graceful Restart for BGP.
`(config-router)# neighbor 10.10.10.2 remote-as 400`	Configure Neighbor for AS-400.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode

(config-router-af)# neighbor 10.10.10.2 activate	Activate neighborship
(config-router-af)#redistribute connected	Redistributing connected Routes inside BGP.
(config-router-af)# neighbor 10.10.10.2 capability graceful-restart	Configure GR capability inside router bgp.
(config-router-af)#end	Exit router BGP mode.

RTR2

#configure terminal	Enter Configuration mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)# ip address 2.2.2.2/32 secondary	Configure ip address on loopback.
(config-if)#exit	Exit interface mode.
(config)#interface xe4	Enter interface mode for eth1.
(config-if)#ip address 10.10.10.2/24	Configure ip address on eth1.
(config-if)#exit	Exit interface mode for eth1.
(config)#interface xe8	Enter interface mode for eth2.
(config-if)#ip address 20.20.20.1/24	Configure ip address on eth2.
(config-if)#exit	Exit interface mode for eth2.
(config)#router bgp 400	Enter router BGP mode.
(config-router)# bgp router-id 2.2.2.22.2.2.2	Configure bgp router-id same as loopback ip address.
(config-router)# bgp graceful-restart	Configure Graceful Restart for BGP.
(config-router)# neighbor 10.10.10.1 remote-as 100	Configure Neighbor for AS-100.
(config-router)# neighbor 20.20.20.2 remote-as 300	Configure Neighbor for AS-300.
(config-router)#address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute connected	Redistributing connected Routes inside BGP.
(config-router-af)# neighbor 10.10.10.1 activate	Activate neighbor
(config-router-af)# neighbor 20.20.20.2 activate	Activate neighbor
(config-router-af)# neighbor 10.10.10.1 capability graceful-restart	Configure GR capability inside router bgp.
(config-router-af)#end	Exit router ospf mode.

RTR3

#configure terminal	Enter Configuration mode.
(config)#interface lo	Enter interface mode for loopback.
(config-if)#ip address 3.3.3.3/32 secondary	Configure ip address on loopback.
(config-if)#exit	Exit interface mode.
(config)#interface xe8	Enter interface mode for eth2.
(config-if)#ip address 20.20.20.2/24	Configure ip address on eth2.

(config-if)#exit	Exit interface mode for eth2.
(config)# router bgp 300	Enter router BGP mode.
(config-router)# bgp router-id 3.3.3.33.3.3.3	Configure bgp router-id same as loopback ip address.
(config-router)# bgp graceful-restart	Configure Graceful Restart for BGP.
(config-router)# neighbor 20.20.20.1 remote-as 400	Configure Neighbor for AS-400.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute connected	Redistributing connected Routes inside BGP.
(config-router-af)# neighbor 20.20.20.1 activate	Activate the neighbor.
(config-router-af)# neighbor 20.20.20.1 capability graceful-restart	Configure GR capability inside router bgp.
(config-router-af)#end	Exit router BGP mode.

Validation

RTR1

```
RTR1#show  bgp neighbors
BGP neighbor is 10.10.10.2, remote AS 400, local AS 100, external link
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 00:03:31
  Last read 00:00:15, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 13 messages, 1 notifications, 0 in queue
  Sent 13 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Graceful restart: advertised, received
     Forwarding states are being preserved
  Community attribute sent to this neighbor (both)
  3 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 10.10.10.1, Local port: 179
Foreign host: 10.10.10.2, Foreign port: 60024
Nexthop: 10.10.10.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:03:36, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

```
#show ip route databaseCodes: K - kernel, C - connected, S - static, R - RIP,
B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C*>1.1.1.11.1.1.1/32 is directly connected, lo, 00:10:23
B*>2.2.2.22.2.2.2/32 [20/0] via 10.10.10.2, xe4, 00:03:56
B*>3.3.3.33.3.3.3/32 [20/0] via 10.10.10.2, xe4, 00:00:56
C*>10.10.10.0/24 is directly connected, xe4, 00:09:37
B*>20.20.20.0/24 [20/0] via 10.10.10.2, xe4, 00:03:56
C*>127.0.0.0/8 is directly connected, lo, 00:28:58

Gateway of last resort is not set
```

RTR2

```
#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked
> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
B*>1.1.1.11.1.1.1/32 [20/0] via 10.10.10.1, xe4, 00:03:52
C*>2.2.2.22.2.2.2/32 is directly connected, lo, 00:07:36
B*>3.3.3.33.3.3.3/32 [20/0] via 20.20.20.2, xe8, 00:00:57

C*>10.10.10.0/24 is directly connected, xe4, 00:07:12 C*>20.20.20.0/24 is
directly connected, xe8, 00:06:31
C*>127.0.0.0/8 is directly connected, lo, 00:25:32

Gateway of last resort is not set

RTR2#show bgp neighbors
BGP neighbor is 10.10.10.1, remote AS 100, local AS 400, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:04:28
  Last read 00:00:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 43 messages, 1 notifications, 0 in queue
  Sent 41 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
```

```
    Graceful restart: advertised, received
     Forwarding states are being preserved
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  3 announced prefixes

 Connections established 4; dropped 3
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 10.10.10.2, Local port: 60050
Foreign host: 10.10.10.1, Foreign port: 179
Nexthop: 10.10.10.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:04:33, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

BGP neighbor is 20.20.20.2, remote AS 300, local AS 400, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 00:04:22
  Last read 00:00:03, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 31 messages, 2 notifications, 0 in queue
  Sent 40 messages, 3 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 2, Offset 0, Mask 0x4
  AF-dependant capabilities:
    Graceful restart: advertised, received
     Forwarding states are being preserved
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  4 announced prefixes

 Connections established 4; dropped 3
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 20.20.20.1, Local port: 179
Foreign host: 20.20.20.2, Foreign port: 56342
Nexthop: 20.20.20.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:04:22, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

RTR2#
```

RTR3

```
#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
B*> 1.1.1.11.1.1.1/32 [20/0] via 20.20.20.1, xe8, 00:01:15
B*> 2.2.2.22.2.2.2/32 [20/0] via 20.20.20.1, xe8, 00:01:15
C*> 3.3.3.3/32 is directly connected, lo
B*> 10.10.10.0/24 [20/0] via 20.20.20.1, xe8, 00:01:15
C*> 20.20.20.0/24 is directly connected, xe8
C*> 127.0.0.0/8 is directly connected, lo

Gateway of last resort is not set
RTR3#show bgp neighbors
BGP neighbor is 20.20.20.1, remote AS 400, local AS 300, external link
  BGP version 4, local router ID 3.3.3.3, remote router ID 2.2.2.2
  BGP state = Established, up for 00:06:47
  Last read 00:00:13, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 45 messages, 1 notifications, 0 in queue
  Sent 38 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Graceful restart: advertised, received
     Forwarding states are being preserved
  Community attribute sent to this neighbor (both)
  4 accepted prefixes
  0 announced prefixes

 Connections established 4; dropped 3
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 20.20.20.2, Local port: 56342
Foreign host: 20.20.20.1, Foreign port: 179
Nexthop: 20.20.20.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:06:52, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

RTR3#
```

Validation After BGP Graceful Restart

RTR2

```
#write
Building configuration... [OK]

#restart bgp graceful
2003 Sep 19 07:20:00.947 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]:
Neighbour
[10.10.10.1] Session down as GR configured/unconfigured
2003 Sep 19 07:20:00.947 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]:
Neighbour
[10.10.10.1] Session down due to config deletion
2003 Sep 19 07:20:00.947 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]:
Neighbour
[20.20.20.2] Session down as GR configured/unconfigured
2003 Sep 19 07:20:00.947 : NOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]:
Neighbour
[20.20.20.2] Session down due to config deletion

#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN, v - vrf leaked

> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
B*>p 1.1.1.1/32 [20/0] via 10.10.10.1, xe4, 00:19:31
C*>2.2.2.2/32 is directly connected, lo, 00:50:45
B*>p 3.3.3.3/32 [20/0] via 20.20.20.2, xe8, 00:19:32
C*>10.10.10.0/24 is directly connected, xe4, 00:50:21 C*>20.20.20.0/24 is
directly connected, xe8, 00:49:40
C*>127.0.0.0/8 is directly connected, lo, 01:08:41 Gateway of last resort is
not set
#show rib forwarding-timer
Protocol-Name GR-State Time Remaining (sec)Disconnected-time BGPACTIVE572001/
06/07 19:50:38
```

RTR1

```
#show ip bgp
BGP table version is 8, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - Labeled, S Stale
Origin codes: i - IGP, e - EGP,? - incomplete

NetworkNext HopMetricLocPrfWeightPath
*>1.1.1.1/32
0.0.0.0010032768?
S>2.2.2.2/32
10.10.10.201000400?
S>3.3.3.3/32
10.10.10.201000400300?
*>10.10.10.0/240.0.0.0010032768?
```

```
S10.10.10.201000400?
S>20.20.20.0/2410.10.10.201000400?

Total number of prefixes 5
```

RTR3

```
#sh ip bgp
BGP table version is 14, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

NetworkNext HopMetricLocPrfWeightPath
S>1.1.1.1/32
20.20.20.101000400 100?
S>2.2.2.2/32
20.20.20.101000400 ?
*>3.3.3.3/32
0.0.0.0010032768?
S>10.10.10.0/2420.20.20.101000400 ?
*>20.20.20.0/240.0.0.0010032768?
S20.20.20.101000400 ?

Total number of prefixes 5
```

CHAPTER 3 BGP Additional Path

Overview

The Border Gateway Protocol (BGP) ADDPATH allows the advertisement of multiple paths through the same peer session for a given prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces the severity of a network failure, thereby improving the control plane convergence in case of network failures.

Feature Characteristics

The advertisement of multiple paths in BGP is made possible by sending a BGP OPEN message to the neighbor with a BGP capability code of 69, which identifies the BGP ADD-PATH capability.

Feature	Characteristics
Address Family Identifier (AFI)	2 octets
Subsequent Address Family Identifier (SAFI)	1 octet
Send/Receive	1 octet

For a given <AFI, SAFI>, the send/receive field in the BGP TLV indicates, the sender is able to:

- Receive multiple paths from its peer (value 1).
- Send multiple paths to its peer (value 2)
- Receive and send multiple paths to its peer (value 3)
- Each alternate path is identified by a Path Identifier in addition to the address prefix

Feature	Characteristics
Path Identifier	4 octets
Length	1 octet
Prefix	variable

Benefits

This feature enables BGP add-path in the vrf address-family. In the event of a next-hop failure, BGP Add-Path improves the BGP control plane convergence time.

Prerequisites

Before configuring BGP additional paths ensure be sure of the following:

- The supported OcNOS router running a compatible release.
- Provide access to the management interface of the router.

- Understand BGP well enough to enable it BGP on an interface.

Configuration

The following sessions displays the detailed information about bgp additional paths topology, configurations, and validations.

Topology

The following topology visually represents how BGP additional paths are configured.

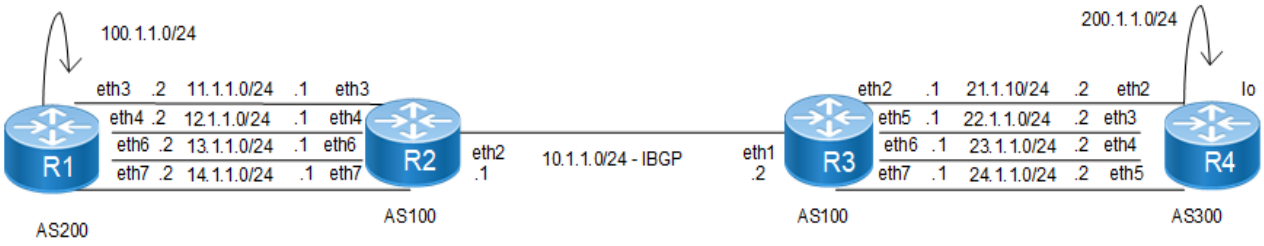

Figure 3-1: BGP Additional Path

R1

Here is the detailed configuration of router R1.

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth2`	Enter the Interface mode to configure interface eth2
`R1(config-if)#ipv6 address 1001::1/64`	Configure an IPv6 address for Interface eth2
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth3`	Enter the Interface mode to configure for Interface eth3
`R1(config-if)#ipv6 address 1002::1/64`	Configure an IPv6 address for Interface eth3
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth4`	Enter the Interface mode to configure Interface eth4
`R1(config-if)#ipv6 address 1003::1/64`	Configure an IPv6 address for interface eth4
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth5`	Enter the Interface mode to configure Interface eth5
`R1(config-if)#ipv6 address 1004::1/64`	Configure an IPv6 address for Interface eth5
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface lo`	Enter Interface mode for loopback lo
`R1(config-if)#ipv6 address 1090::1/64`	Configure IPv6 address for Loopback interface lo
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#router bgp 200`	Enter the Router BGP mode
`R1(config-router)#neighbor 1001::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.

`R1(config-router)#neighbor 1002::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#neighbor 1003::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#neighbor 1004::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#address-family ipv6 unicast`	Enter address-family mode for the neighbor router session to activate.
`R1(config-router-af)#neighbor 1001::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1002::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1003::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1004::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#network 1090::/64`	Activate the neighbor router with a peer address.
`R1(config-router-af)#exit-address-family`	Exit the Address Family mode and return to Router mode.
`R1(config-router)#exit`	Exit the Router BGP mode and enter the Configure mode
`R1(config)#commit`	Apply commit
`R1(config)#exit`	Exit the Configure mode

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter Interface mode for interface eth1
`R2(config-if)#ipv6 address 3001::1/64`	Configure IPv6 address for the interface eth1
`R2(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R2(config-if)#exit`	Exit the Interface mode
`R2(config)#interface eth2`	Enter Interface mode for interface eth2
`R2(config-if)#ipv6 address 1001::2/64`	Configure IPv6 address for the Interface eth2
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#interface eth3`	Enter interface mode for Interface eth3
`R2(config-if)#ipv6 address 1002::2/64`	Configure IPv6 address for the Interface eth3
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#interface eth4`	Enter interface mode for Interface eth4
`R2(config-if)#ipv6 address 1003::2/64`	Configure IPv6 address for the Interface eth4
`R2(config-if)#exit`	Exit the Interface mode
`R2(config)#interface eth5`	Enter interface mode for Interface eth5
`R2(config-if)#ipv6 address 1004::2/64`	Configure IPv6 address for the interface eth5
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#router bgp 100`	Enter the router bgp mode

`R2(config-router)#neighbor 3001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1001::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1002::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1003::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1004::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 1001::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1002::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1003::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1004::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 3001::2 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#exit-address-family`	Exit address family mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#router ipv6 ospf`	Enter Router OSPFv3 mode.
`R2(config-router)#redistribute connected`	Configure Redistribution of Connected networks into OSPF
`R2(config-router)#exit`	Exit the router OSPF mode and enter the configure mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#interface eth1`	Enter Interface mode for Interface eth1
`R3(config-if)#ipv6 address 3001::2/64`	Configure IPv6 address for the Interface eth1
`R3(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth2`	Enter interface mode for Interface eth2
`R3(config-if)#ipv6 address 2001::2/64`	Configure IPv6 address for the Interface eth2
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth3`	Enter Interface mode for the Interface eth3
`R3(config-if)#ipv6 address 2002::2/64`	Configure IPv6 address for the Interface eth3
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth4`	Enter Interface mode for the Interface eth4

R3(config-if)#ipv6 address 2003::2/64	Configure IPv6 address for the Interface eth4
R3(config-if)#exit	Exit the interface mode
R3(config)#interface eth5	Enter Interface mode for Interface eth5
R3(config-if)#ipv6 address 2004::2/64	Configure IPv6 address for the Interface eth5
R3(config-if)#exit	Exit the Interface mode
R3(config)#router bgp 100	Enter the router bgp mode
R3(config-router)#neighbor 3001::1 remote-as 100	Specify a neighbor router with peer address and remote-as for BGP peering.
R3(config-router)#neighbor 2001::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
R3(config-router)#neighbor 2002::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
R3(config-router)#neighbor 2003::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
R3(config-router)#neighbor 2004::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
R3(config-router)#address-family ipv6 unicast	Enter address-family mode for neighbor router session to activate.
R3(config-router-af)#neighbor 2001::1 activate	Activate the neighbor router with peer address.
R3(config-router-af)#neighbor 2002::1 activate	Activate the neighbor router with peer address.
R3(config-router-af)#neighbor 2003::1 activate	Activate the neighbor router with peer address.
R3(config-router-af)#neighbor 2004::1 activate	Activate the neighbor router with peer address.
R3(config-router-af)#neighbor 3001::1 activate	Activate the neighbor router with peer address.
R3(config-router-af)#exit-address-family	Exit address family mode.
R3(config-router)#exit	Exit Router BGP mode
R3(config)#router ipv6 ospf	Enter Router OSPFv3 mode.
R3(config-router)#redistribute connected	Configure Redistribution of Connected networks into OSPF
R3(config-router)#exit	Exit the router OSPF mode and enter theconfigure mode
R3(config)#commit	Apply commit
R3(config)#exit	Exit the configure mode

R4

Here is the detailed configuration of router R4.

R4#configure terminal	Enter the Configure mode.
R4(config)#interface eth2	Enter interface mode for the Interface eth2
R4(config-if)#ipv6 address 2001::1/64	Configure IPv6 address for the Interface eth2
R4(config-if)#exit	Exit the Interface mode
R4(config)#interface eth3	Enter interface mode for Interface eth3
R4(config-if)#ipv6 address 2002::1/64	Configure IPv6 address for the Interface eth3

`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface eth4`	Enter Interface mode for the Interface eth4
`R4(config-if)#ipv6 address 2003::1/64`	Configure IPv6 address for the Interface eth4
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface eth5`	Enter Interface mode for the Interface eth5
`R4(config-if)#ipv6 address 2004::1/64`	Configure IPv6 address for the Interface eth5
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface lo`	Enter interface mode for loopback lo
`R4(config-if)#ipv6 address 9999::1/64`	Configure IPv6 address for Loopback Interface lo
`R4(config-if)#exit`	Exit the interface mode
`R4(config)#router bgp 300`	Enter the router BGP mode
`R4(config-router)#neighbor 2001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2002::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2003::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2004::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R4(config-router-af)#neighbor 2001::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2002::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2003::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2004::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#network 9999::/64`	Activate the neighbor router with peer address.
`R4(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R4(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R4(config)#commit`	Commit the transaction
`R4(config)#exit`	Exit the configure mode

Additional Paths at the Global Level

In the following sessions additional paths at the global level is illustrated.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode

`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to all iBGP neighbors
`R2(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to the Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Commit the transaction
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#bgp additional-paths receive`	Configure R3 to receive additional paths from all iBGP neighbors
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Commit the transaction
`R3(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router.

```
#show bgp ipv6 neighbors 3001::2
BGP neighbor is 3001::2, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.92
  BGP state = Established, up for 00:14:55
  Last read 00:14:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 536 messages, 50 notifications, 0 in queue
  Sent 611 messages, 3 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
```

```
 For address family: IPv6 Unicast
  BGP table version 38, neighbor version 38
  Index 5, Offset 0, Mask 0x20
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  4 announced prefixes

 Connections established 3; dropped 2
Local host: 3001::1, Local port: 38451
Foreign host: 3001::2, Foreign port: 179
Nexthop: 10.12.5.93
Nexthop global: 3001::1
Nexthop local: fe80::5054:ff:fe19:1758
BGP connection: shared network
Last Reset: 00:15:00, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp ipv6 summary
BGP router identifier 10.12.5.93, local AS number 100
BGP table version is 38
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1001::1                    4   200  517        532      38       0       0
04:13:51                1
1002::1                    4   200  520        533      38       0       0
04:13:51                1
1003::1                    4   200  519        532      38       0       0
04:13:51                1
1004::1                    4   200  518        532      38       0       0
04:13:51                1
3001::2                    4   100  588        616      38       0       0
00:15:42                1

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
BGP table version is 38, local router ID is 10.12.5.93
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>   1090::/64        1001::1(fe80::5054:ff:fe9c:b7e6)
                                          0        100          0       200 i
*                     1002::1(fe80::5054:ff:fe0d:f5e)
                                          0        100          0       200 i
*                     1003::1(fe80::5054:ff:fec7:1940)
                                          0        100          0       200 i
```

```
*                     1004::1(fe80::5054:ff:fe62:70d8)
                                        0        100          0       200 i
*>i  9999::/64        2001::1             0        100          0       300
i

Total number of prefixes 2

#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 neighbors 3001::1
BGP neighbor is 3001::1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.93
  BGP state = Established, up for 00:29:37
  Last read 00:29:37, hold time is 90, keepalive interval is 30 seconds
```

```
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 518 messages, 2 notifications, 0 in queue
  Sent 520 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 268, neighbor version 268
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Add-Path Send Capability : received
    Add-Path Receive Capability : advertised
  Community attribute sent to this neighbor (both)
  4 accepted prefixes
  1 announced prefixes

 Connections established 4; dropped 3
Local host: 3001::2, Local port: 179
Foreign host: 3001::1, Foreign port: 38451
Nexthop: 10.12.5.92
Nexthop global: 3001::2
Nexthop local: fe80::5054:ff:fe5d:bb79
BGP connection: shared network
Last Reset: 00:29:37, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
#show bgp ipv6 summary
BGP router identifier 10.12.5.92, local AS number 100
BGP table version is 268
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
2001::1                   4   300  533        537     268      0      0
04:16:42               1
2002::1                   4   300  533        536     268      0      0
04:16:42               1
2003::1                   4   300  537        538     268      0      0
04:16:42               1
2004::1                   4   300  520        521     268      0      0
04:16:38               1
3001::1                   4   100  520        521     268      0      0
00:29:41               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
```

```
BGP table version is 268, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1                0        100          0       200
i
* i                   1004::1                0        100          0       200
i
* i                   1003::1                0        100          0       200
i
* i                   1002::1                0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                            0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                            0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                            0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                            0        100          0       300 i

Total number of prefixes 2

R3#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:08:51 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017
```

© 2025 IP Infusion Inc. Proprietary

Additional Paths Send and Receive at Address-family level

The following session displays the additional paths Send and Receive at Address-family level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send-receive`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`R2(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#bgp additional-paths send-receive`	Configure R3 to send additional paths to and receive additional paths from all the iBGP neighbors
`R3(config-router-af)#bgp additional-paths select all`	Configure R3 to select all available paths to send to all iBGP neighbors
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Apply commit
`R3(config)#exit`	Exit the configure mode

Additional Paths at the Neighbor Level

The following session displays the additional paths at the neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode

`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send-receive`	Configure R2 to send-receive additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#neighbor 3001::1 additional-paths send-receive`	Configure R3 to receive additional paths from the iBGP neighbor R2
`R3(config-router-af)#neighbor 3001::1 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Apply commit
`R3(config)#exit`	Exit the configure mode

Validation

The following validation for router R2 and R3 is shown below.

R2

The following validation is for router R2.

```
#show bgp ipv6 neighbors 3001::2
BGP neighbor is 3001::2, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.92
  BGP state = Established, up for 00:00:29
  Last read 00:00:29, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 588 messages, 51 notifications, 0 in queue
  Sent 664 messages, 4 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
```

```
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 64, neighbor version 64
  Index 5, Offset 0, Mask 0x20
  AF-dependant capabilities:
    Add-Path Send Capability : advertised and received
    Add-Path Receive Capability : advertised and received
  Community attribute sent to this neighbor (both)
  4 accepted prefixes
  4 announced prefixes

 Connections established 5; dropped 4
Local host: 3001::1, Local port: 179
Foreign host: 3001::2, Foreign port: 39326
Nexthop: 10.12.5.93
Nexthop global: 3001::1
Nexthop local: fe80::5054:ff:fe19:1758
BGP connection: shared network
Last Reset: 00:00:29, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp ipv6 summary
BGP router identifier 10.12.5.93, local AS number 100
BGP table version is 64
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1001::1                   4   200  561        578      64       0       0
04:35:32               1
1002::1                   4   200  564        579      64       0       0
04:35:32               1
1003::1                   4   200  563        578      64       0       0
04:35:32               1
1004::1                   4   200  562        578      64       0       0
04:35:32               1
3001::2                   4   100  640        669      64       0       0
00:00:35               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
BGP table version is 64, local router ID is 10.12.5.93
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>   1090::/64        1001::1(fe80::5054:ff:fe9c:b7e6)
                                          0        100          0       200 i
*                     1002::1(fe80::5054:ff:fe0d:f5e)
```

```
                                          0        100          0       200 i
*                     1003::1(fe80::5054:ff:fec7:1940)
                                          0        100          0       200 i
*                     1004::1(fe80::5054:ff:fe62:70d8)
                                          0        100          0       200 i
*>i  9999::/64        2001::1               0        100          0       300
i
* i                   2002::1               0        100          0       300
i
* i                   2003::1               0        100          0       300
i
* i                   2004::1               0        100          0       300
i

Total number of prefixes 2

#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

```
#show bgp ipv6 9999::/64
BGP routing table entry for 9999::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1001::1 1002::1 1003::1 1004::1
  300
    2001::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:39 2017

  300
    2002::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017

  300
    2003::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: 2
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017

  300
    2004::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: 3
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017
```

R3

The following validation is for router R3.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:39 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
```

```
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: 2
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: 3
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

R3#show bgp ipv6 9999::/64
BGP routing table entry for 9999::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2002::1 2003::1 2004::1
  300
    2001::1(fe80::5054:ff:fe46:f549) from 2001::1 (10.12.5.90)
    (fe80::5054:ff:fe46:f549)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:32 2017

  300
    2002::1(fe80::5054:ff:fed2:4666) from 2002::1 (10.12.5.90)
    (fe80::5054:ff:fed2:4666)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:27 2017

  300
    2003::1(fe80::5054:ff:fe0d:b565) from 2003::1 (10.12.5.90)
    (fe80::5054:ff:fe0d:b565)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:37 2017

  300
    2004::1(fe80::5054:ff:feb5:9a71) from 2004::1 (10.12.5.90)
    (fe80::5054:ff:feb5:9a71)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:44 2017

#show bgp ipv6
BGP table version is 283, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
```

```
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1                0        100          0       200
i
* i                   1002::1                0        100          0       200
i
* i                   1003::1                0        100          0       200
i
* i                   1004::1                0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                            0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                            0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                            0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                            0        100          0       300 i

Total number of prefixes 2

#show bgp ipv6 summary
BGP router identifier 10.12.5.92, local AS number 100
BGP table version is 283
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
2001::1                   4   300  556        562     282      0      0
04:28:07               1
2002::1                   4   300  556        560     283      0      0
04:28:07               1
2003::1                   4   300  560        563     282      0      0
04:28:07               1
2004::1                   4   300  543        546     283      0      0
04:28:03               1
3001::1                   4   100  551        553     283      0      0
00:04:18               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6 neighbors 3001::1
BGP neighbor is 3001::1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.93
  BGP state = Established, up for 00:05:02
  Last read 00:05:02, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 550 messages, 3 notifications, 0 in queue
  Sent 553 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
```

```
 Minimum time between advertisement runs is 5 seconds
For address family: IPv4 Unicast
 BGP table version 1, neighbor version 1
 Index 5, Offset 0, Mask 0x20
 Community attribute sent to this neighbor (both)
 0 accepted prefixes
 0 announced prefixes

For address family: IPv6 Unicast
 BGP table version 283, neighbor version 283
 Index 1, Offset 0, Mask 0x2
 AF-dependant capabilities:
   Add-Path Send Capability : advertised and received
   Add-Path Receive Capability : advertised and received
 Community attribute sent to this neighbor (both)
 4 accepted prefixes
 4 announced prefixes

 Connections established 6; dropped 5
Local host: 3001::2, Local port: 39326
Foreign host: 3001::1, Foreign port: 179
Nexthop: 10.12.5.92
Nexthop global: 3001::2
Nexthop local: fe80::5054:ff:fe5d:bb79
BGP connection: shared network
Last Reset: 00:05:07, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

Selection of all Additional Paths at the Address-family Level

The following are the configurations and validations for additional paths at the address-family level.

R2

Here is the detailed configuration of router R2.

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Enter BGP router mode
`(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`(config-router-af)#bgp additional-paths send-receive`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`(config)#commit`	Apply the commit
`(config)#exit`	Exit the configure mode

Selection of all Additional Paths at the Neighbor Level

The following are the configurations and validations for additional paths at the neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017
```

```
  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:01 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017
```

Selection of Best 2 Additional Paths at AF Level

The following are the configurations and validations for best 2 additional paths at AF level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#bgp additional-paths select best 2`	Configure R2 to select best 2 out of all available paths to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Selection of Best 2 Additional Paths at the Neighbor Level

The following are the configurations and validations for best 2 additional paths at neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths best 2`	Configure R2 to advertise best 2 out of all available paths to R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::
BGP routing table entry for 1090::/64
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
```

```
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
#show bgp ipv6
BGP table version is 407, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1               0        100          0       200
i
* i                   1002::1               0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                           0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                           0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                           0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                           0        100          0       300 i

Total number of prefixes 2
```

Selection of Best 3 Additional Paths at the AF Level

The following are the configurations and validations for best 3 additional paths at the AF level.

R2

THere is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#bgp additional-paths select best 3`	Configure R2 to select best 3 out of all available paths to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Selection of Best 3 Additional Paths at the Neighbor Level

The following are the configurations and validations for best 3 additional paths at neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths best 3`	Configure R2 to advertise best 3 out of all available paths to R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017
```

```
  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:11 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:53 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:53 2017

#show bgp ipv6
BGP table version is 410, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1               0        100          0       200
i
```

```
* i                    1003::1                0        100          0       200
i
* i                    1002::1                0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                            0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                            0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                            0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                            0        100          0       300 i

Total number of prefixes 2
```

Implementation Examples

The following examples shows the bgp additional path entries.

```
OcNOS#show ip bg summary
BGP router identifier 33.33.33.1, local AS number 400
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1.1.1.2                  4   100   19         17       2      0      0
00:06:35               0
Total number of neighbors 1
Total number of Established sessions 1
```

CLI Commands

The BGP additional path introduces the following configuration commands.

bgp additional-paths send-receive

Use this command is to enable BGP additional paths send and receive global mode commands in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path send and receive global mode commands in ipv4 vrf address-family.

Command Syntax

```
bgp additional-paths (send-receive)
```

```
no bgp additional-paths (send-receive)
```

Parameters

`send`	Send additional paths to neighbors.
`receive`	Receive additional paths from neighbors.
`send-receive`	Send and receive additional paths from neighbors.

`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

bgp additional-paths select best 3

Use this command is to enable BGP additional best 3 paths in global mode ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path select best 3 paths in global mode ipv4 vrf address-family.

Command Syntax

```
bgp additional-paths(select best 3)
```

```
no bgp additional-paths(select best 3)
```

Parameters

`best`	Select best N paths.
`3`	Number of best paths in additional paths to be selected.
`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
```

```
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

bgp additional-paths select all

Use this command to enable BGP additional paths select all in global mode commands in ipv4 vrf address-family.

Use the `no` parameter with this command to disable all selected BGP add-paths in global mode commands in ipv4 vrf address-family..

Command Syntax

```
bgp additional-paths(select all)
```

```
no bgp additional-paths(select all)
```

Parameters

`all`	Select all available paths.
`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

neighbor A.B.C.D additional-paths send | receive | send-receive

Use this command to enable BGP add-path at neighbor level to send and receive neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level to send and receive neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths (send|receive|send-receive|)
```

```
no neighbor A.B.C.D additional-paths (send|receive|send-receive|)
```

Parameters

`send`	Send additional paths to neighbors.

`receive`	Receive additional paths from neighbors.
`send-receive`	Send and receive additional paths from neighbors.

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

neighbor A.B.C.D additional-paths all

Use this command to enable BGP add-path at all neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths all
```

```
no neighbor A.B.C.D additional-paths all
```

Parameters

`all`	Select all available paths

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

neighbor A.B.C.D additional-paths best <2-3>

Use this command to enable BGP add-path at all neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths all
no neighbor A.B.C.D additional-paths all
```

Parameters

`best`	Select best N paths.
`<2-3>`	Number of best paths in additional paths to be selected.

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

Troubleshooting

BGP additional paths for DC is a new feature when it is configured neighbour level resets only the particular peer and global level resets all the peers.

Abbreviations

List key terms used in this document and add the term and explanation to our existing Glossary.

Acronym	Description
BGP	Border Gateway Protocol

CLI	Command Line Interface
TLV	Type Length Values

CHAPTER 4 Hide the Remote AS using the neighbor local-as Command

Overview

In a network, an Autonomous System (AS) is available to define a set of IP routing prefixes that are under a common administration policy control. These defined routing policies are used by other connected routers o n the Internet. When an AS is configured in Border Gateway Protocol (BGP), it is used to share routing information to connected peers. The `neighbor local-as` CLI command configures the AS number to be used with External Border Gateway Protocol (EBGP) peers. By default, the configured AS number is included in the AS-PATH message that is exchanged between the peers.

When a BGP router, configured in one network, connects to another router on the network, it will automatically share routing information with the AS number of both the local and remote routers in the AS-PATH message with other connected, external peers. For example, if a router ISP1-R, accesses services from another router, ISP2-R, ISP1-R router will share routing information with local and remote AS numbers in the AS-PATH message when services are merged. This allows the external peers to learn the AS numbers of remote routers not connected to it (in this case, the AS number of ISP2-R). It is not desirable to disclose the AS number of remote routers to external peers.

To avoid advertising the remote peer's AS number, OcNOS provides an option in the `neighbor local-as` CLI to not include (`no-prepend`) the remote AS number and replace (`replace-as`) it with alternate AS number. Configuring an alternate AS in the BGP neighbor system, provides the ability to hide the AS number of the remote router that actually shares the services. Thus, the AS number of the BGP router that is actually providing services is unknown to the external peer.

Hence, the existing `neighbor local-as` CLI command has been modified in this release.

Feature Characteristics

The `neighbor local-as` CLI is enhanced to hide and replace the AS number of the remote routers not connected to external peer. Two new options '`no-prepend`' and '`replace-as`' have been added. These options replace the AS number with an alternate AS number in the AS_PATH and BGP OPEN message. Hence, the AS of the remote router is unknown to the respective neighbor peer.

Benefits

The actual Autonomous System number is never shared to the external network.

Configuration

The following configuration assumes the router R1 and R2 is assigned with AS300 and AS100 respectively.

Topology

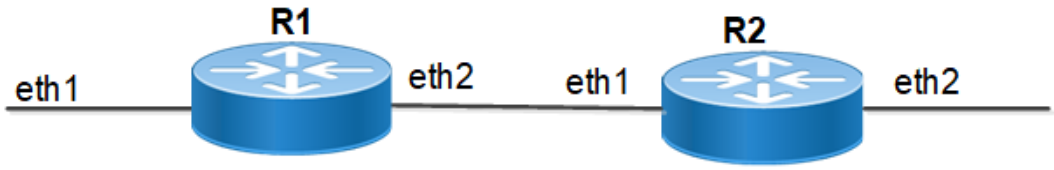

Figure 4-2: Disparate Autonomous System Number

R1

Perform the following configuration on R1 router.

`#configure terminal`	Enter configure mode.
`R1(config)#router bgp 300`	Start the BGP process with the Autonomous System number 300
`R1(config-router)#neighbor 10.10.10.2 remote-as 200`	Establish BGP session with neighbor that has AS number 200
`R1(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`R1(config-router-af)#neighbor 10.10.10.2 activate`	Enable the neighbor 10.10.10.2 router to exchange address family routes
`R1(config-router-af)#redistribute connected`	Redistribute information from connected routes
`R1(config-router-af)#exit-address-family`	Exit address-family IPv4 unicast mode
`R1(config-router)#commit`	Commit the configurations

R2

Perform the following configuration on R2 router.

`#configure terminal`	Enter configure mode
`R2(config)#router bgp 100`	Start the BGP process with the Autonomous System number 100
`R2(config-router)#neighbor 10.10.10.1 remote-as 300`	Establish BGP session with neighbor 10.10.10.1 that has AS number 300
`R2(config-router)#neighbor 10.10.10.1 local-as 200 no-prepend replace-as`	Replace the AS number 300 with AS number 200 that should be used with the neighbor 10.10.10.1
`R2(config-router)#address-family ipv4 unicast`	Enable the neighboring router to exchange address family routes
`R2(config-router-af)#neighbor 10.10.10.2 activate`	Enable the neighbor 10.10.10.2 router to exchange address family routes
`R2(config-router-af)#redistribute connected`	Redistribute information from the connected routes
`R2(config-router-af)#exit-address-family`	Exit address-family ipv4 unicast mode
`R2(config-router)#commit`	Commit the configurations

Validation

Check the AS number 300 running on R1. It has established a BGP connection with 10.10.10.2 router that has AS number of 200.

R1

```
OcNOS#show running-config bgp
!
router bgp 300
 neighbor 10.10.10.2 remote-as 200
 !
 address-family ipv4 unicast
 redistribute connected
 redistribute static
 neighbor 10.10.10.2 activate
 exit-address-family
!
OcNOS#
OcNOS#show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 300
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
10.10.10.2               4   200  185        181       3      0       0  00:00:28       2

Total number of neighbors 1

Total number of Established sessions 1
OcNOS#

OcNOS#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, ce1, 1d14h18m
B            30.30.30.0/24 [20/0] via 10.10.10.2, ce1, 00:00:18
C            40.40.40.0/24 is directly connected, xe33, 1d13h40m
C            127.0.0.0/8 is directly connected, lo, 1d14h23m
Gateway of last resort is not set
OcNOS#
```

Check if the AS number 100 for R2 has been replaced with AS number 200 before sharing the information with R1.

R2

```
OcNOS#show running-config bgp
!
router bgp 100
 neighbor 10.10.10.1 remote-as 300
 neighbor 10.10.10.1 local-as 200
 !
 address-family ipv4 unicast
 redistribute connected
 redistribute static
 neighbor 10.10.10.1 activate
 exit-address-family
!
OcNOS#
OcNOS#show ip bgp summary
BGP router identifier 10.10.10.2, local AS number 100
BGP table version is 2
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
10.10.10.1              4   300  180       186      2      0      0       00:00:39       2

Total number of neighbors 1

Total number of Established sessions 1
```

Check if the AS number for R2 is changed to 100 and R1 shares AS 100 in the AS-PATH message.

R1

```
OcNOS#
OcNOS#
OcNOS#show ip bgp
BGP table version is 4, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric    LocPrf   Weight Path
*> 10.10.10.0/24    0.0.0.0             0         100      32768  ?
*                   10.10.10.2          0         100       0     200 100 ?
*> 30.30.30.0/24    10.10.10.2          0         100       0     200 100 ?
*> 40.40.40.0/24    0.0.0.0             0         100       32768  ?

Total number of prefixes 3
```

Revised CLI Commands

neighbor local-as

The `neighbor local-as` CLI is enhanced to hide and replace the AS number of the remote routers not connected to external peer. Two new options '`no-prepend`' and '`replace-as`' have been added. These options replace the AS number with an alternate AS number in the AS_PATH and BGP OPEN message. Hence, the AS of the remote router is unknown to the respective neighbor peer.

For the complete command reference, refer to neighbor local-as CLI in BGP Commands section.

Abbreviations

Acronym	Description
ASN	Autonomous System Number
EBGP	External Border Gateway Protocol

CHAPTER 5 TCP MSS configuration for BGP neighbors

Overview

The manual configuration between the routing devices establishes the BGP peer that creates a TCP session.

This feature enables the configuration of TCP Maximum Segment Size (MSS) that defines the maximum segment size in a single TCP segment during a communication session. TCP segment is a unit of data transmitted in a TCP connection. TCP uses three-way handshake process for initial establishment of a TCP connection. In the three-way handshake process, the sending host sends a SYN packet. Once the receiving host receives the SYN packet, it acknowledges and sends back a SYN-ACK packet to the sending host. Once the sending host receives the SYN-ACK packet from the receiving host, it sends an ACK packet, establishing a reliable connection. In this three way handshake process, the MSS is negotiated between the BGP neighbors.

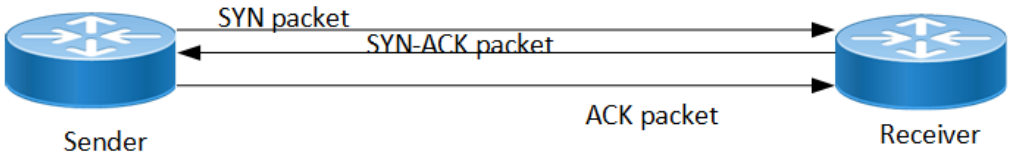

Figure 5-1: Three-way handshake

Feature Characteristics

The configuration of the TCP MSS for BGP neighbors helps the neighbors adjust the MSS value of the TCP SYN packet. Configure the TCP MSS through the CLI and NetConf interface. The configurable MSS range is offered from 40-1440 bytes. By default, the MTU value for ethernet cable is 1500 bytes. When configuring the highest MSS value that is 1440, the total MSS becomes 1440 bytes (MSS) plus 20 bytes (IP Header Size), 20 bytes (TCP Header), and Ethernet header which does not cross the default path MTU value.

Figure 5-2: TCP MSS for BGP neighbor

Benefits

By default, the interface MTU value determines the MSS value of a packet. When the interface MTU value exceeds the default ethernet path MTU value of 1500 bytes, the MSS value also crosses the default ethernet path MTU value, resulting in packet fragmentation. The configuration of the specific MSS value limits the packet size irrespective of the interface MTU value, preventing packet fragmentation.

Prerequisites

Requires the knowledge on TCP handshake and BGP neighbor discovery.

Configuration

This section shows the procedure to configure TCP MSS between BGP peers.

Topology

The below example shows the configuration required to enable BGP on an interface. PE1 and RR1 are routers belonging to the same Autonomous System (AS) with the Autonomous System Number (ASN) as AS100, connecting to network 10.1.1.0/24. First, define the routing process and the ASN to which the routers belong. Then, define BGP neighbors to start exchanging routing updates and configure the TCP MSS for BGP between PE1 and RR1 devices.

Figure 5-3: Device topology for BGP

Configuration

The configuration shows how to configure the TCP MSS value for the BGP peer.

PE1

`PE1#configure terminal`	Enter Configuration mode.
`PE1(config)#interface lo`	Enter interface mode for loopback.
`PE1(config-if)#ip address 1.1.1.1/32 secondary`	Specify the interface IP address `1.1.1.1`.
`PE1(config-if)#exit`	Exit the interface mode.
`PE1(config)#interface xe1`	Enter interface mode for `xe1`.
`PE1(config-if)#ip address 10.1.1.1/24`	Specify the IP address `10.1.1.1` for the interface.
`PE1(config-if)#exit`	Exit interface mode for `xe1`.
`PE1(config)#router bgp 100`	Define the routing process. The number 100 specifies the ASN of PE1.
`PE1(config-router)#bgp router-id 1.1.1.1`	Configure bgp router-id same as loopback IP address`1.1.1.1`.
`PE1(config-router)#neighbor 10.1.1.2 remoteas 100`	Define BGP neighbors, and establish a TCP session. `10.1.1.2` is the IP address of the neighbor and 100 is the neighbor's ASN.

`PE1(config-router)#neighbor 10.1.1.2 tcp-mss 800`	Configure TCP MSS value.
`PE1(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode.
`PE1(config-router-af)#neighbor 10.1.1.2 activate`	Activate neighbor with IP address `10.1.1.2` in the IPv4 address family.
`PE1(config-router-af)#redistribute connected`	Redistributing connected routes inside BGP.
`PE1(config-router-af)#exit-address-family`	Exit address-family mode.
`PE1(config-router)#commit`	Commit the candidate configuration to the running configuration.

RR1

`RR1#configure terminal`	Enter configuration mode.
`RR1(config)#interface lo`	Enter interface mode for loopback.
`RR1(config-if)#ip address 2.2.2.2/32 secondary`	Specify the interface address `2.2.2.2`.
`RR1(config-if)#exit`	Exit interface mode.
`RR1(config)#interface xe47`	Enter interface mode for `xe47`.
`RR1(config-if)#ip address 10.1.1.2/24`	Specify IP address `10.1.1.2/24` for the interface.
`RR1(config-if)#exit`	Exit interface mode for `xe47`.
`RR1(config)#router bgp 100`	Define the routing process. The number 100 specifies the ASN of RR1.
`RR1(config-router)#bgp router-id 2.2.2.2`	Configure BGP router-id same as loopback IP address `2.2.2.2`.
`RR1(config-router)#neighbor 10.1.1.1 remoteas 100`	Define BGP neighbors, and establish a TCP session. `10.1.1.1` is the ip address of the neighbor and 100 is the neighbor's ASN.
`RR1(config-router)#neighbor 10.1.1.1 passive`	Configure BGP neighbor `10.1.1.1` passive.
`RR1(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode
`RR1(config-router-af)#neighbor 10.1.1.1 activate`	Activate the neighbor in the IPv4 address family.
`RR1(config-router-af)#neighbor 10.1.1.1 route-reflector-client`	Configure RR1 as the Route-Reflector (RR) and neighbor PE1 as its client.
`RR1(config-router-af)#redistribute connected`	Redistributing connected routes inside BGP.
`RR1(config-router-af)#exit-address-family`	Exit address-family mode.
`RR1(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

PE1

```
PE1#show bgp summary
```

```
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd
10.1.1.2                 4   100  171        170       1      0      0  00:00:11
               0

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#sh bgp neighbors
BGP neighbor is 10.1.1.2, remote AS 100, local AS 100, internal link, peer index
: 2
  BGP version 4, local router ID 10.1.1.1, remote router ID 10.1.1.2
  BGP state = Established, up for 00:07:29
  Last read 00:00:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 43 messages, 1 notifications, 0 in queue
  Sent 46 messages, 4 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 10.1.1.1, Local port: 34738
Foreign host: 10.1.1.2, Foreign port: 179
TCP MSS: (800), Advertise TCP MSS: (800), Send TCP MSS: (800),  Receive TCP MSS:
 (536)
Sock FD : (25)
Nexthop: 10.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:08:45, due to Administratively Reset (Cease Notification sent)
```

RR1

```
RR1#show bgp summary
BGP router identifier 2.2..2.2, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd
10.1.1.1                 4   100    2          3       1      0       0  00:00:26
               0

Total number of neighbors 1

Total number of Established sessions 1

RR1#sh bgp neighbors
BGP neighbor is 10.1.1.1, remote AS 100, local AS 100, internal link, peer index
: 2
  BGP version 4, local router ID 10.1.1.2, remote router ID 10.1.1.1
  BGP state = Established, up for 00:08:31
  Last read 00:00:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 46 messages, 4 notifications, 0 in queue
  Sent 47 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 10.1.1.2, Local port: 179
Foreign host: 10.1.1.1, Foreign port: 34738
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (800),  Receive TCP MSS:
(536)
Sock FD : (22)
Nexthop: 10.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:09:52, due to BGP Notification received
```

New CLI Commands

neighbor tcp-mss

Use this command to set the BGP TCP MSS of a neighbor.

Use the `no` parameter with this command to remove a TCP MSS setting from a BGP neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) tcp-mss <40-1440>
no neighbor (A.B.C.D|X:X::X:X|WORD) tcp-mss
```

For BGP unnumbered mode:

```
neighbor WORD tcp-mss <40-1440>
no neighbor WORD tcp-mss
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group range command. When you specify this parameter, the command applies to all peers in the group.
`<40-1440>`	Configure TCP MSS

Default

By default, `neighbor tcp-mss` is disabled.

Command Mode

Router mode, address family-vrf mode and BGP unnumbered mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.72 tcp-mss 1000
(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 tcp-mss 900
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 tcp-mss 800
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
ACK	Acknowledgment
BGP	Border Gateway Protocol
TCP	Transmission Control Protocol
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
SYN	Synchronize

Glossary

The following provides definitions for key terms used throughout this document.

BGP	BGP is an exterior gateway protocol to exchange route information and interconnect various networks on the global internet.
BGP neighbor	BGP neighbors, called peers, are established by manual configuration among routers to create a TCP session on port 179, which exchanges routing information between two systems, defined by their Autonomous System Numbers (ASNs).
MSS	MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted. TCP - TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices.
TCP	TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices.
TCP segment	TCP segment is a unit of data transmitted in a TCP connection. The segment consists of header and payload. The header contains the control information to manage the transmission, and the payload contains the actual data that needs to be transmitted.

CHAPTER 6 BGP Bogon Prefix Filtering IPv4

Overview

Bogon filters block invalid or reserved IP addresses, known as Bogon prefixes, from propagating through a network. These prefixes typically include IP ranges not allocated by IANA and should not appear in the global routing table. In BGP, bogon filtering prevents the acceptance of such invalid routes. By default in OcNOS, BGP does not filter bogon prefixes in incoming route updates, except for Class-E Experimental IPv4 Addresses (240.0.0.0/4) and Multicast Addresses (224.0.0.0/4 and 224.0.0.0/24). The new CLI command bgp enable-bogon-filtering allows users to manage this filtering process, adding flexibility to network security.

Feature Characteristics

A list of BOGON prefixes in OcNOS that should not route on the internet:

- **0.0.0.0/8:** Reserves addresses for special use, typically for traffic from hosts without an assigned IP address.
- **127.0.0.0/8:** Reserves addresses for loopback, used for internal host communication.
- **192.0.0.0/24:** Reserves addresses for IETF Protocol Assignments.
- **192.0.2.0/24:** Reserves addresses for documentation and examples.

BOGON and Class-E Experimental Address Range

240.0.0.0/4: Class-E addresses, reserved for experimental purposes. Avoid using them in production environments.

Note:

- After filtering or allowing prefixes, BGP bogon filtering policy changes apply only to new incoming updates. The system does not automatically refresh previously filtered or allowed prefixes. When making changes to the BGP bogon filtering configuration, it is recommended to perform a hard reset of BGP or reboot the device.
- No filtering applies to IPv6 prefixes, as BGP allows all IPv6 prefixes by default.

Benefits

Security: Enabling bogon prefix filtering improves network security by preventing the propagation of invalid or reserved IP addresses.

Flexibility: The option to enable or disable filtering provides network administrators control over BGP route updates.

Compliance: Ensures adherence to best practices by allowing for the exclusion of non-routable IP address ranges.

Prerequisites

Based on the running configuration provided for R1 and R2, configure BGP between two routers (R1 and R2) with the same Autonomous System (AS) number and IPv4 unicast address family. This configuration allows both routers to establish a BGP session and exchange routes.

```
R1#show running-config
!
hostname R1
```

```
!
ip vrf management
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe27
 ip address 100.1.1.1/24
!
router bgp 100
 neighbor 100.1.1.2 remote-as 100
 !
 address-family ipv4 unicast
 redistribute static
 neighbor 100.1.1.2 activate
 exit-address-family
 !
 exit
!
!
end

R2#show running-config
!
hostname R2
!
ip vrf management
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
```

```
!
interface xe27
 ip address 100.1.1.2/24
!
router bgp 100
 neighbor 100.1.1.1 remote-as 100
 !
 address-family ipv4 unicast
 neighbor 100.1.1.1 activate
 exit-address-family
 !
 exit
!
!
end
```

Configuration

Ensure the OcNOS routers properly filter out bogon prefixes and enhance network security by preventing the routing of invalid or reserved IP addresses with the below configurations

Topology

Two routers (R1 and R2) connect via interface `xe27`, using the same AS number (`AS 100`), to establish a BGP session and exchange routes within the IPv4 unicast address family.

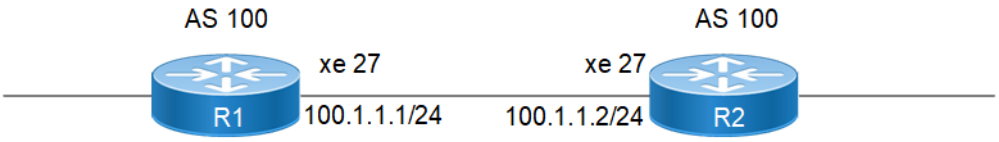

BGP Bogon Filter

Note: Before configuration, meet all Prerequisites.

Configure Invalid IP Addresses on R1 and R2

- Configure the invalid (bogon) routes before enabling BGP bogon filtering:

```
R2(config)#ip route 12.0.0.0/8 xe27
R2(config)#ip route 192.0.0.0/24 xe27 << Bogon prefixes
```

- Verify the injected routes (`12.0.0.0/8 and 192.0.0.0/24`) in the BGP table.

```
R2#show ip bgp
BGP table version is 6, local router ID is 100.1.1.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric     LocPrf    Weight Path
*>   12.0.0.0         0.0.0.0                  0        100       32768  ?
*>   192.0.0.0        0.0.0.0                  0        100       32768  ?

Total number of prefixes 2
R2#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
S            12.0.0.0/8 [1/0] is directly connected, xe27, installed 00:07:25, last update 00:07:25 ago
C            100.1.1.0/24 is directly connected, xe27, installed 00:10:25, last update 00:10:25 ago
C            127.0.0.0/8 is directly connected, lo, installed 00:57:51, last update 00:57:51 ago
S            192.0.0.0/24 [1/0] is directly connected, xe27, installed 00:07:25, last update 00:07:25 ago

Gateway of last resort is not set
```

Enable BGP Bogon Filtering on R1 and R2

Enable the bogon filtering command to block invalid prefixes. This command automatically removes invalid or reserved prefixes from the BGP table.

```
R2(config)#bgp enable-bogon-filtering
```

Validation

Verify the `192.0.0.0/24` bogon prefix is removed from the BGP table.

```
R2#show ip bgp
BGP table version is 7, local router ID is 100.1.1.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf   Weight Path
*>   12.0.0.0         0.0.0.0               0        100       32768  ?

Total number of prefixes 1

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
S            12.0.0.0/8 [1/0] is directly connected, xe27, installed 00:07:25, last update 00:07:25 ago
C            100.1.1.0/24 is directly connected, xe27, installed 00:10:25, last update 00:10:25 ago
C            127.0.0.0/8 is directly connected, lo, installed 00:57:51, last update 00:57:51 ago

Gateway of last resort is not set
```

BGP Bogon Prefix Filtering Command

The bogon prefix filtering introduces the following configuration command:

bgp enable-bogon-filtering

Use this command to enable or disable the filtering of bogon prefixes from NLRIs in incoming BGP updates.

Note: After filtering or allowing prefixes, BGP bogon filtering policy changes apply only to new incoming updates. The system does not automatically refresh previously filtered or allowed prefixes. When making changes to the BGP bogon filtering configuration, it is recommended to perform a hard reset of BGP or reboot the device.

Command Syntax

```
bgp enable-bogon-filtering
```

Parameters

None

Default

Disabled.

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.3.

Example

In OcNOS, enable BGP bogon filtering and perform a hard reset or device reboot as recommended.

```
OcNOS#configure terminal
OcNOS(config)#bgp enable-bogon-filtering
OcNOS(config)#commit
2098 Apr 24 21:11:07.665 : OcNOS : BGP : CRITI : % Changing the bogon filter
does not automatically refresh the permitted/filtered routes. A hard reset or
device reboot is recommended.
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Bogon Filter	A mechanism that blocks invalid or reserved IP addresses, known as Bogon prefixes, from entering a network through routing updates.
Bogon Prefix	IP address ranges that are not allocated by IANA and should not appear in the global routing table.
Internet Assigned Numbers Authority (IANA)	The organization responsible for allocating global IP addresses and managing other critical Internet resources.
Autonomous System (AS)	A collection of IP networks managed by a single entity that presents a unified routing policy to the Internet.

Border Gateway Protocol Command Reference

CHAPTER 1 BGP Commands

This chapter describes the BGP configuration commands.

- address-family
- aggregate-address
- auto-summary
- bgp additional-paths select all
- bgp additional-paths send-receive
- bgp additional-paths select best 3
- bgp aggregate-nexthop-check
- bgp always-compare-med
- bgp as-local-count
- bgp bestpath as-path ignore
- bgp bestpath as-path multipath-relax
- bgp bestpath compare-confed-aspath
- bgp bestpath compare-routerid
- bgp bestpath dont-compare-originator-id
- bgp bestpath med
- bgp bestpath tie-break-on-age
- bgp client-to-client reflection
- bgp cluster-id
- bgp confederation vrf identifier
- bgp confederation vrf peers
- bgp dampening
- bgp default local-preference
- bgp deterministic-med
- bgp enforce-first-as
- bgp extended-asn-cap
- bgp fast-external-failover
- bgp graceful-restart
- bgp g-shut
- bgp g-shut-capable
- bgp g-shut-local-preference
- bgp log-neighbor-changes
- bgp nexthop-trigger delay
- bgp nexthop-trigger enable
- bgp rfc1771-path-select
- bgp rfc1771-strict

© 2025 IP Infusion Inc. Proprietary

- bgp router-id
- bgp scan-time
- bgp table-map
- bgp update-delay
- clear bgp (A.B.C.D|X:X::X:X)
- clear bgp *
- clear bgp <1-4294967295>
- clear bgp <1-4294967295> l2vpn evpn
- clear bgp A.B.C.D l2vpn evpn
- clear bgp dampening
- clear bgp external
- clear bgp flap-statistics
- clear bgp peer-group
- clear bgp peer-group WORD l2vpn evpn
- clear bgp statistics
- clear ip bgp * (description LINE|)*
- clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|)
- clear ip bgp A.B.C.D
- clear ip bgp A.B.C.D vrf
- clear ip bgp all vrf (VRFNAME|all|default) (description LINE|)
- clear ip bgp table-map
- clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|)
- debug bgp
- distance bgp
- exit-address-family
- ip as-path access-list
- ip community-list <1-99>
- ip community-list <100-500>
- ip community-list expanded
- ip community-list standard
- ip community-list WORD
- ip extcommunity-list <1-99>
- ip extcommunity-list <100-500>
- ip extcommunity-list expanded
- ip extcommunity-list standard
- l2vpn-unnumbered-mode
- match ip peer
- match large-community
- max-paths

- max-paths eigbp
- neighbor A.B.C.D additional-paths all
- neighbor activate
- neighbor advertisement-interval
- neighbor allowas-in
- neighbor as-origination-interval
- neighbor attribute-unchanged
- neighbor authentication-key
- neighbor A.B.C.D additional-paths best <2-3>
- neighbor A.B.C.D additional-paths send | receive | send-receive
- neighbor capability graceful-restart
- neighbor capability orf prefix-list
- neighbor collide-established
- neighbor de-activate
- neighbor default-originate
- neighbor description
- neighbor disallow-infinite-holdtime
- neighbor distribute-list
- neighbor dont-capability-negotiate
- neighbor ebgp-multihop
- neighbor enforce-multihop
- neighbor extended-optional-param
- neighbor fall-over bfd
- neighbor filter-list
- neighbor g-shut
- neighbor g-shut-timer
- neighbor limit
- neighbor local-as
- neighbor maximum-prefix
- neighbor next-hop-self
- neighbor optional-as
- neighbor override-capability
- neighbor passive
- neighbor peer-group
- neighbor WORD peer-group range
- neighbor port
- neighbor prefix-list
- neighbor remote-as
- neighbor remove-private-AS

© 2025 IP Infusion Inc. Proprietary

- neighbor restart-time
- neighbor route-map
- neighbor route-reflector-client
- neighbor route-server-client
- neighbor send-community
- neighbor send-community large
- neighbor shutdown
- neighbor strict-capability-match
- neighbor tcp-mss
- neighbor timers
- neighbor unsuppress-map
- neighbor update-source
- neighbor version
- neighbor weight
- neighbor WORD peer-group
- network
- network synchronization
- redistribute
- restart bgp gracefull
- router bgp
- set large-community
- snmp restart bgp
- synchronization
- timers bgp
- v4-unnumbered-mode
- bgp unnumbered-mode

address-family

Use the address family command to enter the IPv4 address family mode allowing configuration of address-family specific parameters. To leave the address family mode and return to the Configure mode use the `exit-address-family` command.

This command configures the routing exchange between Provider Edge (PE) and Customer Edge (CE) devices. The BGP sessions between PE routers can carry different types of routes (VPN-IPv4 and IPv4 routes). Address families are used to control the type of BGP session. Configure a BGP address family for each VRF configured on the PE router and a separate address family to carry VPN-IPv4 routes between PE routers. All non VPN BGP neighbors are defined using router mode. All VPN BGP neighbors are defined under its associated address family mode. The BGP process with no address-family specified is the default address-family, where any sessions are configured that either are not associated with a VRF or are used to carry IPv4 routes.

Use the `no` parameter with this command to disable the address-family configurations.

Command Syntax

```
address-family ipv4
address-family ipv4 (unicast|multicast)
address-family ipv4 vrf NAME
address-family rtfilter unicast
address-family vpn4 unicast
no address-family ipv4 vrf NAME
no address-family ipv4 (unicast|multicast)
no address-family rtfilter unicast
```

Parameters

`ipv4`	IPv4 address family
`unicast`	Unicast address prefixes
`multicast`	Multicast address prefixes
`vrf`	Virtual Private Network (VPN) routing/forwarding instance
`NAME`	VPN routing/forwarding instance name
`unicast`	Unicast address prefixes
`rtfilter`	Route target filter: on an iBGP peer or Route Reflector (RR), only send IPv4 and IPv6 prefixes to PE routers when a PE has a VRF that imports those specific prefixes.
`unicast`	Unicast address prefixes
`unicast`	Unicast address prefixes

Applicability

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 7657
(config-router)#neighbor 3ffe:506::1 remote-as 7657
(config-router)#neighbor 3ffe:506::1 interface eth1

#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 3ffe:506::1 activate
(config-router-af)#exit-address-family
```

aggregate-address

Use this command to configure BGP aggregate entries.

Aggregates are used to minimize the size of routing tables. Aggregation combines the characteristics of several different routes and advertises a single route.The `aggregate-address` command creates an aggregate entry in the BGP routing table if any more-specific BGP routes are available in the specified range. Using the `summary-only` parameter advertises the prefix only, suppressing the more-specific routes to all neighbors. In the following example Router1 will propagate network 172.0.0.0 and suppresses the more specific route 172.10.0.0.

The `as-set` parameter creates an aggregate entry advertising the path for this route, consisting of all elements contained in all paths being summarized. Use this parameter to reduce the size of path information by listing the AS number only once, even if it was included in multiple paths that were aggregated. The `as-set` parameter is useful when aggregation of information results in an incomplete path information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
aggregate-address A.B.C.D/M
aggregate-address A.B.C.D/M as-set
aggregate-address A.B.C.D/M as-set summary-only
aggregate-address A.B.C.D/M summary only
aggregate-address A.B.C.D/M summary-only as-set
no aggregate-address A.B.C.D/M
```

Parameters

`A.B.C.D/M`	Aggregate prefix
`as-set`	Generate AS set path information
`summary-only`	Filter more specific routes from updates

Default

By default, aggregate address A.B.C.D/M is disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#aggregate-address 10.0.0.0/8 as-set summary-only

(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#no aggregate-address 10.0.0.0/8
```

auto-summary

Use this command to enable sending summarized routes by a BGP speaker to its peers in the router configuration mode or in the address-family configuration mode. Auto-summary is used by a BGP router to advertise summarized routes to its peers. Auto-summary can be enabled if certain routes have already been advertised: in this case, configuring auto-summary advertises the summarized routes first, then corresponding non-summarized routes are withdrawn. If certain routes have already been advertised, and auto-summary is disabled, non-summarized routes are first advertised, then the corresponding summarized routes are withdrawn from all the connected peers.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
auto-summary
```

```
no auto-summary
```

Parameters

None

Default

By default, auto-summary is disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv4 unicast
(config-router-af)#auto-summary
```

bgp aggregate-nexthop-check

Use this command to set the BGP option to perform aggregation only when next-hop matches the specified IP address.

Use the `no` parameter with this command to disable this functionality.

Command Syntax

```
bgp aggregate-nexthop-check
no bgp aggregate-nexthop-check
```

Parameters

None

Default

By default, bgp aggregate nexthop check is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp aggregate-nexthop-check
```

bgp always-compare-med

Use this command to compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. Multi Exit Discriminator (MED) is used in best path selection by BGP. MED is compared after BGP attributes weight, local preference, AS-path and origin have been compared and are equal. MED comparison is done only among paths from the same autonomous system (AS). Use `bgp always-compare-med` command to allow comparison of MEDs from different ASs. The MED parameter is used to select the best path. A path with lower MED is preferred. If the bgp table shows the following and the always-compare-med is enabled:

```
Route1: as-path 400, med 300
Route2: as-path 200, med 200
Route3: as-path 400, med 250
```

Route1 is compared to Route2. Route2 is best of the two (lower MED). Next, Route2 is compared to Route3 and Route2 is chosen best path again (lower MED). If always-compare-med was disabled, MED is not taken into account when Route1 and Route2 are compared, because of different ASs and MED is compared for only Route1 and Route3. In this case, Route3 would be the best path. The selected route is also affected by the `bgp deterministic-med` command. Please see `bgp deterministic-med` command for details. If this command is used to compare MEDs for all paths, it should be configured on every BGP router in the AS.

Use the `no` parameter with this command to disallow the comparison.

Command Syntax

```
bgp always-compare-med
no bgp always-compare-med
```

Parameters

None

Default

By default, bgp always compare med is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp always-compare-med
```

bgp as-local-count

Use this command to set the number of times the local-AS (Autonomous System) is to be prepended.

Use the `no` parameter with this command to stop prepending the local AS count.

Command Syntax

```
bgp as-local-count <2-64>
no bgp as-local-count <2-64>
```

Parameter

`<2-64>`	The number of times the local-AS is to be prepended

Default

By default, bgp as local count is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp as-local-count 55

(config)#router bgp 100
(config-router)#no bgp as-local-count 55
```

bgp bestpath as-path ignore

Use this command to prevent the router from considering the autonomous system (AS) path length as a factor in the algorithm for choosing a best path route.

Use the `no` parameter with this command to allow the router to consider the AS path length in choosing a best path route.

Command Syntax

```
bgp bestpath as-path ignore
```

```
no bgp bestpath as-path ignore
```

Parameters

None

Default

By default, bgp bestpath as-path ignore is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath as-path ignore

(config)#router bgp 100
(config-router)#no bgp bestpath as-path ignore
```

bgp bestpath as-path multipath-relax

Use this command to relax the “same AS-Path” requirement so any candidate eBGP AS-Path with the same AS-path length might be used for eBGP load-balancing.

Note: This feature does not load-balance between eBGP and iBGP paths.

Normally eBGP load-balancing requires the candidate routes to be equal-cost paths with identical BGP attributes having the same weight, Local-Pref, AS-Path (both the AS numbers and the AS path length), origin, MED, and different next-hop.

Use the `no` parameter with this command to return to normal operation.

Command Syntax

```
bgp bestpath as-path multipath-relax
```

```
no bgp bestpath as-path multipath-relax
```

Parameters

None

Default

By default, as-path multipath-relax is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath as-path multipath-relax

(config)#router bgp 100
(config-router)#no bgp bestpath as-path multipath-relax
```

bgp bestpath compare-confed-aspath

Use this command to allow comparing of the confederation AS path length. This command specifies that the AS confederation path length must be used when available in the BGP best path decision process. It is effective only when bgp bestpath as-path ignore command has not been used.

Use the `no` parameter with this command to ignore consideration of AS confederation path length in BGP best path selection.

Command Syntax

```
bgp bestpath compare-confed-aspath
```

```
no bgp bestpath compare-confed-aspath
```

Parameters

None

Default

BGP receives routes with identical eBGP paths from eBGP peers and selects the first route received as the best path.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath compare-confed-aspath

(config)#router bgp 100
(config-router)#no bgp bestpath compare-confed-aspath
```

bgp bestpath compare-routerid

Use this command to compare router IDs for identical eBGP paths. When comparing similar routes from peers, the BGP router does not consider the router ID of the routes. By default, it selects the first received route. Use this command to include router ID in the selection process; similar routes are compared and the route with the lowest router ID is selected. The router IS is the highest IP address on the router, with preference given to loopback addresses. Router ID can be manually set by using the bgp router-id command.

Use the `no` parameter with this command to disable this functionality.

Command Syntax

```
bgp bestpath compare-routerid
```

```
no bgp bestpath compare-routerid
```

Parameters

None

Default

BGP receives routes with identical eBGP paths from eBGP peers and selects the first route received as the best path.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath compare-routerid

(config)#router bgp 100
(config-router)#no bgp bestpath compare-routerid
```

bgp bestpath dont-compare-originator-id

Use this command to change the default bestpath selection by not comparing an originator-ID for an identical EBGP path.

Use the `no` parameter with this command to disable this functionality.

Command Syntax

```
bgp bestpath dont-compare-originator-id
```

```
no bgp bestpath dont-compare-originator-id
```

Parameters

None

Default

BGP receives routes with identical eBGP paths from eBGP peers and selects the first route received as the best path.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath dont-compare-originator-id

(config)#router bgp 100
(config-router)#no bgp bestpath dont-compare-originator-id
```

bgp bestpath med

Use this command to specify two MED (Multi Exit Discriminator) attributes, `confed` and `missing-as-worst`.

The `confed` attribute enables MED comparison along paths learned from confederation peers. The MEDs are compared only if there is no external Autonomous System (an AS not within the confederation) in the path. If there is an external autonomous system in the path, the MED comparison is not made. For example in the following paths, the MED is not compared with Route3 as it is not in the confederation. MED is compared for Route1 and Route2 only.

```
Path1 = 32000 32004, med=4
Path2 = 32001 32004, med=2
Path3 = 32003 1, med=1
```

The `missing-as-worst` attribute to 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. If `missing-as-worst` is disabled, the missing MED is assigned the value of 0, making the path with the missing MED attribute the best path.

Use the `no` parameter with this command to prevent BGP from considering the MED attribute in comparing paths.

Command Syntax

```
bgp bestpath med confed missing-as-worst
bgp bestpath med (confed|missing-as-worst|remove-recv-med|remove-send-med)
bgp bestpath med missing-as-worst confed
no bgp bestpath med confed missing-as-worst
no bgp bestpath med (confed|missing-as-worst|remove-recv-med|remove-send-med)
no bgp bestpath med missing-as-worst confed
```

Parameters

`confed`	Compare MED along confederation paths
`missing-as-worst`	Treat missing MED as the least preferred one
`remove-recv-med`	Remove received MED attribute
`remove-send-med`	Remove sent MED attribute

Command Mode

Router mode

Default

By default, MED value is zero.

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
```

```
(config)#router bgp 100
(config-router)#bgp bestpath med missing-as-worst

(config)#router bgp 100
(config-router)#bgp bestpath med remove-recv-med
(config-router)#no bgp bestpath med remove-recv-med

(config)#router bgp 100
(config-router)#bgp bestpath med remove-send-med
(config-router)#no bgp bestpath med remove-send-med
```

bgp bestpath tie-break-on-age

Use this command to always select a preferred older route even when the `bgp bestpath compare-routerid` command is configured.

Use the `no` parameter with this command to disable this functionality.

Command Syntax

```
bgp bestpath tie-break-on-age
```

```
no bgp bestpath tie-break-on-age
```

Parameters

None

Default

By default, tie-break-on-age is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp bestpath tie-break-on-age

(config)#router bgp 100
(config-router)#no bgp bestpath tie-break-on-age
```

bgp client-to-client reflection

Use this command to configure routers as route reflectors. Route reflectors are used when all Interior Border Gateway Protocol (iBGP) speakers are not fully meshed. If the clients are fully meshed the route reflector is not required, use `no bgp client-to-client reflection` command to disable the client-to-client route reflection.

Use the `no` parameter with this command to turn off client-to-client reflection.

Command Syntax

```
bgp client-to-client reflection
```

```
no bgp client-to-client reflection
```

Parameters

None

Default

When a router is configured as a route reflector, client-to-client reflection is enabled by default.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp client-to-client reflection

(config)#router bgp 100
(config-router)#no bgp client-to-client reflection
```

bgp cluster-id

Use this command to configure the cluster ID if the BGP cluster has more than one route reflector. A cluster includes route reflectors and its clients. Usually, each cluster is identified by the router ID of its single route reflector but to increase redundancy sometimes a cluster may have more than one route reflector. All router reflectors in such a cluster are then identified by a cluster ID. The `bgp cluster-id` command is used to configure the 4 byte cluster ID for clusters with more than one route reflectors.

Use the `no` parameter with this command (without any arguments) to remove a previously configured route reflector cluster ID.

Command Syntax

```
bgp cluster-id <1-4294967295>
bgp cluster-id A.B.C.D
no bgp cluster-id
```

Parameters

`<1-4294967295>`	Route reflector ID as a 32-bit quantity
`A.B.C.D`	Route reflector ID in an IPv4 address format

Default

By default, cluster id is set bgp cluster id

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following configuration creates a cluster-id 5 including two route-reflector-clients.

```
#configure terminal
(config)#router bgp 100
(config-router)#neighbor 2.2.2.2 remote-as 200
(config-router)#neighbor 3.3.3.3 remote-as 200
(config-router)#neighbor 5.5.5.5 remote-as 200
(config-router)#neighbor 6.6.6.6 remote-as 200
(config-router)#bgp cluster-id 5
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 3.3.3.3 route-reflector-client
(config-router-af)#neighbor 5.5.5.5 route-reflector-client
```

© 2025 IP Infusion Inc. Proprietary

bgp confederation vrf identifier

Use this command to specify a BGP confederation identifier.

Use the `no` parameter with this command to remove a BGP confederation identifier.

Command Syntax

```
bgp confederation (vrf NAME |) identifier <1-4294967295>
no bgp confederation (vrf NAME |) identifier
```

Parameter

`identifier <1-4294967295>`	Routing domain confederation AS number.
`vrf NAME`	Specify the user-defined VRF instance name.

Default

No default value is specified.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced VRF parameter in OcNOS version 6.6.0. This CLI is configured at global mode. The default vrf remains unchanged in global mode.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp confederation vrf vrf1 identifier 1000
```

bgp confederation vrf peers

Use this command to configure the Autonomous Systems (AS) that belong to a confederation. A confederation allows an AS to be divided into several ASs. The AS is given a confederation identifier. External BGP (eBGP) routers view only the whole confederation as one AS. Each AS is fully meshed within itself and is visible internally to the confederation.

Use the `no` parameter with this command to remove an autonomous system from the confederation.

Command Syntax

```
bgp confederation (vrf NAME |) peers <1-4294967295>
no bgp confederation (vrf NAME |) peers <1-4294967295>
```

Parameter

`<1-4294967295>`	AS numbers of eBGP peers that are in the same confederation
`vrf NAME`	Specify the user-defined VRF instance name.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced VRF parameter in OcNOS version 6.6.0. This CLI is configured at global mode. The default vrf remains unchanged in global mode.

Examples

In the following configuration example, the neighbor 172.210.30.2 and 172.210.20.1 have iBGP connection within AS 100, neighbor 173.213.30.1 is a BGP connection with a confederation peer 200 and neighbor 6.6.6.6 has an eBGP connection to external AS 300.

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp confederation vrf vrf1 identifier 300
(config-router)#bgp confederation vrf vrf1 peers 200
(config-router)#neighbor 172.210.30.2 remote-as 100
(config-router)#neighbor 172.210.20.1 remote-as 100
(config-router)#neighbor 173.213.30.1 remote-as 200
(config-router)#neighbor 6.6.6.6 remote-as 300
```

In this configuration, the neighbor 5.5.5.4 has an eBGP connection to confederation 300.

```
#configure terminal
(config)#router bgp 500
(config-router)#neighbor 5.5.5.4 remote-as 300
```

bgp dampening

Use this command to enable BGP route dampening and set various parameters. Route dampening minimizes the instability caused by route flapping. A penalty is added for every flap in a flapping route. As soon as the total penalty reaches the `suppress` limit the advertisement of the route is suppressed. This penalty is decayed according to the configured `half time` value. Once the penalty is lower than the `reuse` limit, the route advertisement is un-suppressed. The dampening information is purged from the router once the penalty becomes less than half of the `reuse` limit.

Use the `no` parameter with this command to unset BGP dampening parameters.

Command Syntax

```
bgp dampening
bgp dampening <1-45>
bgp dampening <1-45> <1-20000> <1-20000> <1-255>
bgp dampening <1-45> <1-20000> <1-20000> <1-255> <1-45>
bgp dampening route-map WORD
no bgp dampening
```

Parameters

`<1-45>`	Reachability half-life time for the penalty in minutes. The time for the penalty to decrease to one-half of its current value.
`<1-20000>`	Value to start reusing a route. When the penalty for a suppressed route decays below the reuse value, the routes become unsuppressed.
`<1-20000>`	Value to start suppressing a route. When the penalty for a route exceeds the suppress value, the route is suppressed
`<1-255>`	Maximum duration to suppress a stable route in minutes.
`<1-45>`	Un-reachability half-life time for the penalty in minutes.
`route-map`	Route map to specify criteria for dampening.
`WORD`	Route-map name.

Defaults

The default reachability half-life is 15 minutes.

The default reuse limit is 750.

The default suppress limit is 2000.

The default max-suppress value is 4 times the half-life time, or 60 minutes.

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp dampening 20 800 2500 80 25
```

bgp default local-preference

Use this command to change the default local preference value.Local preference indicates the preferred path when there are multiple paths to the same destination. The path having a higher preference is preferred. The preference is sent to all routers and access servers in the local autonomous system.

Use the `no` parameter with this command to revert to the default value for local preference.

Command Syntax

```
bgp default local-preference <0-4294967295>
no bgp default local-preference
no bgp default local-preference <0-4294967295>
```

Parameter

`<0-4294967295>` Local preference value

Default

By default, local preference value is 100

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp default local-preference 2345555
```

bgp deterministic-med

Use this command to compare the Multi Exit Discriminator (MED) variable when choosing among routes advertised by different peers in the same autonomous system. MED is compared after BGP attributes weight, local preference, AS-path and origin have been compared and are equal.

For a correct comparison result, enable this command on all routers in a local AS. After enabling this command, all paths for the same prefix are grouped together and arranged according to their MED value. Based on this comparison, the best path is then chosen. This command compares MED variable when choosing routes advertised by different peers in the same AS, to compare MED, when choosing routes from neighbors in different ASs use the bgp always-compare-med command.

When the bgp deterministic-med command is enabled, routes from the same AS are grouped together, and the best routes of each group are compared. If the BGP table showed:

```
Route1: as-path 200, med 300, internal
Route2: as-path 400, med 200, internal
Route3: as-path 400, med 250, external
```

BGP would have a group of Route1 and a second group of Route2 and Route3 (the same ASs). The best of each group is compared. Route1 is the best of its group because it is the only route from AS 200. Route1 is compared to the Route2, the best of group AS 400 (the lower MED). Since the two routes are not from the same AS, the MED is not considered in the comparison. The external BGP route is preferred over the internal BGP route, making Route3 the best route; the preferred route would be different if `always-compare-med` command is enabled (See `always-compare-med` command).

Use the `no` parameter with this command to disallow this setting.

Command Syntax

```
bgp deterministic-med
```

```
no bgp deterministic-med
```

Parameters

None

Default

By default, bgp deterministic med is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp deterministic-med

(config)#router bgp 100
(config-router)#no bgp deterministic-med
```

bgp enforce-first-as

Use this command to enforce the first AS for eBGP routes. This command specifies that any updates received from an external neighbor that do not have the neighbor's configured Autonomous System (AS) at the beginning of the AS_PATH in the received update must be denied. Enabling this feature adds to the security of the BGP network by not allowing traffic from unauthorized systems.

Using the `no` parameter with this command to disable this feature.

Command Syntax

```
bgp enforce-first-as

no bgp enforce-first-as
```

Parameters

None

Default

By default, enforce-first-as is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp enforce-first-as

(config)#router bgp 100
(config-router)#no bgp enforce-first-as
```

bgp extended-asn-cap

Use this command to configure a BGP router to send 4-octet ASN capabilities. If attempting to change the AS capability from 2 to 4 or 4 to 2, a prompt occurs to remove the VRF configuration (if it exists), and reconfiguration is required, because the route distinguisher (RD) configuration would have been created with the current (2 octet or 4 octet) capability, and must be reconfigured before attempting to change the capability.

While loading from a saved configuration with AS4 capability and BGP VRF configuration, the capability will not be changed because of the above described reason.

Use the `no` parameter with this command to prevent a BGP router from sending 4-octet ASN capabilities.

Command Syntax

```
bgp extended-asn-cap
```

```
no bgp extended-asn-cap
```

Parameters

None

Default

By default, the bgp extended ASN capability and Four-octet capabilities are disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp extended-asn-cap
```

bgp fast-external-failover

Use this command to reset a BGP session immediately, if the interface used for BGP connection goes down.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
bgp fast-external-failover
no bgp fast-external-failover
```

Parameters

None

Default

By default, fast-external failover is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp fast-external-failover
```

bgp graceful-restart

Use this command to enable BGP graceful-restart capabilities. The restart-time parameter is used for setting the maximum time that a graceful-restart neighbor waits to come back up after a restart. This value is applied to all neighbors unless you explicitly override it by configuring the corresponding value on the neighbor. The stalepath-time parameter is used to set the maximum time to preserve stale paths from a gracefully restarted neighbor. All stalepaths, unless reinstated by the neighbor after a re-establishment, will be deleted at the expiration of this timer.

Use the `no` parameter with this command to restore the router to its default state.

Command Syntax

```
bgp graceful-restart
bgp graceful-restart graceful-reset
bgp graceful-restart restart-time <1-3600>
bgp graceful-restart stalepath-time <1-3600>
no bgp graceful-restart
no bgp graceful-restart graceful-reset
no bgp graceful-restart restart-time
no bgp graceful-restart stalepath-time
```

Parameters

`graceful-reset`	The BGP daemon is not restarted, so that any changes in network configurations that cause BGP reset do not affect packet forwarding.
`restart-time`	Maximum time needed for neighbors to restart. Default is 90 seconds.
`<1-3600>`	Delay value in seconds.
`stalepath-time`	Maximum time to retain stale paths from restarting neighbors. Default is 360 seconds.
`<1-3600>`	Delay value in seconds.

Default

By default, the maximum time for neighbors to restart is 90 seconds.

By default, the maximum time to retain stale paths from restarting neighbors is 360 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#bgp graceful-restart

#configure terminal
(config)#router bgp 10
```

```
(config-router)#no bgp graceful-restart
```

bgp g-shut

Use this command to gracefully shut down all BGP IPv4 sessions under this router. The BGP graceful shutdown feature reduces packet loss during maintenance activity.

Use the `no` parameter with this command to bring up all the sessions under this router after completion of maintenance activity using the `bgp g-shut` command.

Command Syntax

```
bgp g-shut
```

```
no bgp g-shut
```

Parameters

None

Default

By default, bgp g-shut is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp g-shut

#configure terminal
(config)#router bgp 100
(config-router)#no bgp g-shut
```

bgp g-shut-capable

Use this command to enable the graceful shutdown capability at the router level and make available the graceful-shutdown related commands at the router and BGP neighbor levels.

Use the `no` parameter with this command to disable the graceful shutdown capability on a router.

Note: The graceful shutdown capability cannot be disabled on a router that is in a graceful shutdown state until it comes out this state--after the graceful shutdown has been initiated and the impacted BGP sessions are up again.

Command Syntax

```
bgp g-shut-capable
```

```
no bgp g-shut-capable
```

Parameters

None

Default

By default, the graceful shutdown capability is disabled at the router level

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp g-shut-capable

#configure terminal
(config)#router bgp 100
(config-router)#no bgp g-shut-capable
```

bgp g-shut-local-preference

Use this command to sets the local preference of the router to use during graceful shutdown. The local preference value indicates the preferred path when there are multiple paths to the same destination in a single routing database. The path with a higher preference value is the preferred one. The preferred path is sent to all routers and access servers in the local autonomous system.

Use the `no` parameter with this command to revert to the default setting.

Note: The bgp g-shut-local-preference command relies on the bgp g-shut-capable configuration. When user remove the bgp g-shut-capable command from the configuration, it automatically deletes the bgp g-shut-local-preference setting as well. This dependency means that the bgp g-shut-local-preference cannot function or remain configured without the presence of the bgp g-shut-capable command.

Command Syntax

```
bgp g-shut-local-preference <0-4294967295>
no bgp g-shut-local-preference
```

Parameters

`<0-4294967295>` Local preference value

Default

By default, the local preference value is set to 0

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp g-shut-local-preference 22
```

bgp log-neighbor-changes

Use this command to enable logging of status change messages without turning on debug bgp commands. OcNOS has many logging services for neighbor status, including `debug bgp fsm` and `debug bgp events`. However, these commands cause system performance degradation. If you need to log neighbor status changes only, IP Infusion Inc. recommends turning off all debug commands and using the `bgp log-neighbor-changes` command instead. A sample output of the log is:

```
%Protocol-Severity-Events: Message-text
```

A sample output of the log for an `interface down` event is:

```
%BGP-5-ADJCHANGE: neighbor 10.10.0.24 Down Interface flap
```

This command logs the following events:

- BGP Notification Received
- Erroneous BGP Update Received
- User reset request
- Peer time-out
- Peer Closing down the session
- Interface flap
- Router ID changed
- Neighbor deleted
- Member added to peer group
- Administrative shutdown
- Remote AS changed
- RR client configuration modification
- Soft reconfiguration modification

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
bgp log-neighbor-changes
no bgp log-neighbor-changes
```

Parameters

None

Default

By default, bgp log neighbor changes is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#router bgp 100
(config-router)#bgp log-neighbor-changes
```

bgp nexthop-trigger delay

Use this command to set the delay time for nexthop address tracking. This command configures the delay interval between routing table walks for nexthop delay tracking, after which BGP does a routing table scan on receiving a nexthop change trigger from NSM. The time period determines how long BGP waits before it walks the full BGP table to determine which prefixes are affected by the nexthop changes, after it receives the trigger from NSM about one or more nexthop changes.

Use the `no` parameter with this command to reset the timer value to the default value.

Command Syntax

```
bgp nexthop-trigger delay <1-100>
no bgp nexthop-trigger delay
```

Parameter

`<1-100>`	Nexthop trigger delay interval in seconds

Default

By default, nexthop-trigger delay time is 5 seconds

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp nexthop-trigger delay 6

#configure terminal
(config)#no bgp nexthop-trigger delay
```

bgp nexthop-trigger enable

Use this command to enable nexthop address tracking. Nexthop address tracking is an event-driven notification system that monitors the status of routes installed in the Routing Information Base (RIB) and reports nexthop changes that affect internal BGP (iBGP) or external BGP (eBGP) prefixes directly to the BGP process. This improves the overall BGP convergence time, by allowing BGP to respond rapidly to nexthop changes for routes installed in the RIB.

If nexthop tracking is enabled after certain routes are learned, the registration of all nexthops for selected BGP routes is done after the nexthop tracking feature is enabled. If nexthop tracking is disabled, and if there are still some selected BGP routes, BGP de-registers the nexthops of all selected BGP routes from NSM.

Use the `no` parameter with this command to disable this feature. If the `no` command is given when nexthop tracking is in the process of execution, an error appears and nexthop tracking is not disabled. However, if the nexthop tracking timer is running at the time of negation, the nexthop tracking timer is stopped, and nexthop tracking is disabled.

Command Syntax

```
bgp nexthop-trigger enable
```

```
no bgp nexthop-trigger enable
```

Parameters

None

Default

By default, nexthop address tracking is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp nexthop-trigger enable
```

bgp rfc1771-path-select

Use this command to set RFC 1771 compatible path selection.

Use the `no` parameter with this command to revert this setting.

Command Syntax

```
bgp rfc1771-path-select
no bgp rfc1771-path-select
```

Parameters

None

Default

Standard compatible path selection mechanism.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp rfc1771-path-select
```

bgp rfc1771-strict

Use this command to set the origin path attribute to “IGP” when the origin is a protocol such as RIP, OSPF, or ISIS as specified in RFC 1771. Otherwise, the origin is always set to “incomplete” which is the industry standard.

Use the `no` parameter with this command to revert this setting.

Command Syntax

```
bgp rfc1771-strict
no bgp rfc1771-strict
```

Parameters

None

Default

By default, bgp rfc1771 strict is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bgp rfc1771-strict
```

bgp router-id

Use this command to manually configure a fixed router ID as a BGP router identifier. When this command is used to configure a fixed router ID, the current router identifier is overridden and the peers are reset.

Use the `no` parameter with this command to remove a manually configured fixed router ID.

Command Syntax

```
bgp router-id A.B.C.D
no bgp router-id
no bgp router-id A.B.C.D
```

Parameter

`A.B.C.D`	Router ID in an IPv4 address format

Default

Once the BGP router-id is elected, it may be re-elected in the following cases:

a. When an explicit BGP router-id is configured/un-configured

b. When the router's (global) router-id is set/unset (holds true when (a) is not applicable),

c. When the BGP process is cleared (holds true when both (a) & (b) are not applicable and the IP address(es) on the active interfaces are updated, which may result in a change in the router's router-id).

If no loopback interface is configured, the highest IP address is the BGP router-id.

When a loopback interface is configured, the BGP router-id is set to the IP address of the loopback interface.

Note: IP Infusion Inc. recommends that you always configure a router identifier to avoid unpredictable behavior if the address of a loopback interface changes.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp router-id 10.1.2.7

(config)#router bgp 100
(config-router)#no bgp router-id 10.1.2.7
```

bgp scan-time

Use this command to configure scanning intervals of BGP routers. This interval is the period after which router checks the validity of the routes in its database. To disable BGP scanning, set the scan-time interval to 0 seconds.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
bgp scan-time <0-60>
no bgp scan-time
no bgp scan-time <0-60>
```

Parameter

`<0-60>`	Scanning interval in seconds

Default

By default, scan-time interval is 60 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp scan-time 10
```

bgp table-map

Use this command to enable or disable suppression/modification of incoming BGP updates to IP RIB/FIB table installation.

In a dedicated route reflector, all the routes it receives may not be required to be stored or only few selected routes need to be stored, because it may not lie in the data path.

Table maps are particularly useful to attain this restriction. Table-map command can be used in two ways:

- When a simple table-map command is given, the route map referenced in the table-map command shall be used to set certain properties (such as the traffic index) of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.
- When the option ‘filter’ is given in the table map command, the route map referenced is used to control whether a BGP route is to be downloaded to the IP RIB (hence the filter). A BGP route is not downloaded to the RIB if it is denied by the route map.

Use this command in Address Family mode to set the table map rule per an IPv4 or IPv6 family.

Use the `no` parameter with this command to remove the table-map rule.

Command Syntax

```
table-map WORD (|filter)
```

Parameter

`WORD`	Specify the route-map name to apply.
`filter`	Filer the routes. If present, the incoming routes are pruned as per the rule specified in route-map-name. If not, it is used to alter the incoming packet information.

Default

All BGP routes will be downloaded to IP RIB

Command Mode

Address Family IPv4 mode, and Address Family IPv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows how to set the table-map command without filter for BGP for all address families.

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#table-map abc
```

The following example shows how to set the table-map command with filter for BGP for all address families.

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#table-map abc filter
```

The following example shows how to set the table-map command without filter for BGP for an IPv6 address family.

```
(config)#router bgp 100
(config-router)#address-family ipv6
(config-router-af)#table-map abc
```

The following example shows how to set the table-map command with filter for BGP for an IPv6 address family.

```
(config)#router bgp 100
(config-router)#address-family ipv6
(config-router-af)#table-map abc filter
```

bgp update-delay

Use this command to set the update delay for a graceful-restart capable router. The update-delay value is the maximum time a graceful-restart capable router, which is restarting, will defer route-selection and advertisements to all its graceful-restart capable neighbors. This maximum time starts from the instance the first neighbor attains established state after restart. The restarting router prematurely terminates this timer when end-of-RIB markers are received from all its graceful-restart capable neighbors.

Use the `no` form of this command to set to the update delay to its default value.

Command Syntax

```
bgp update-delay <1-3600>
no bgp update-delay
no bgp update-delay <1-3600>
```

Parameters

`<1-3600>`	Delay interval in seconds

Default

By default, update-delay value is 120 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#bgp update-delay 345
```

clear bgp (A.B.C.D|X:X::X:X)

Use this command to reset a BGP neighbor address.

Command Syntax

```
clear bgp (A.B.C.D|X:X::X:X|WORD)
clear bgp (A.B.C.D|X:X::X:X) in
clear bgp (A.B.C.D|X:X::X:X) in prefix-filter
clear bgp (A.B.C.D|X:X::X:X) out
```

Parameters

`A.B.C.D`	IPv4 neighbor address.
`X:X::X:X`	IPv6 neighbor address.
`WORD`	Interface name
`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp 3.3.3.3
```

clear bgp *

Use this command to reset the BGP connection for all peers.

Command Syntax

```
clear bgp *
clear bgp * in
clear bgp * in prefix-filter
clear bgp * out
clear bgp * soft
clear bgp * soft in
clear bgp * soft out
clear ip bgp *
clear ip bgp * in
clear ip bgp * in prefix-filter
clear ip bgp * out
clear ip bgp * soft
clear ip bgp * soft in
clear ip bgp * soft out
clear ip bgp * ipv4 (unicast|multicast) in
clear ip bgp * ipv4 (unicast|multicast) in prefix-filter
clear ip bgp * ipv4 (unicast|multicast) out
clear ip bgp * ipv4 (unicast|multicast) soft
clear ip bgp * ipv4 (unicast|multicast) soft in
clear ip bgp * ipv4 (unicast|multicast) soft out
```

Parameters

`in`	Incoming advertised routes should be cleared.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes of all the address families.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.
`ipv4`	Clear incoming advertised routes.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
`in`	Clear incoming advertised routes.
`prefix-filter`	

	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear bgp *
#clear ip bgp * ipv4 unicast in prefix-filter
```

clear bgp <1-4294967295>

Use this command to reset a BGP connection for all peers in a specified Autonomous System.

Command Syntax

```
clear bgp <1-4294967295>
clear bgp <1-4294967295> in
clear bgp <1-4294967295> in prefix-filter
clear bgp <1-4294967295> out
clear bgp <1-4294967295> soft
clear bgp <1-4294967295> soft in
clear bgp <1-4294967295> soft out
clear ip bgp <1-4294967295>
clear ip bgp <1-4294967295> in
clear ip bgp <1-4294967295> in prefix-filter
clear ip bgp <1-4294967295> out
clear ip bgp <1-4294967295> soft
clear ip bgp <1-4294967295> soft in
clear ip bgp <1-4294967295> soft out
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) in
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) in prefix-filter
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) out
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) soft
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) soft in
clear ip bgp <1-4294967295> ipv4 (unicast|multicast) soft out
```

Parameters

`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.
`ipv4`	Clear incoming advertised routes.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
	Clear incoming advertised routes.
`in`	Clear incoming advertised routes.

`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear bgp 4294967277
#clear ip bgp 4294967277
```

clear bgp <1-4294967295> l2vpn evpn

Use this command to reset the session for neighbors with a specific Autonomous System Number (ASN) for L2VPN EVPN.

Command Syntax

```
clear bgp <1-4294967295> l2vpn evpn soft (in | out | )
```

Parameters

`bgp <1-4294967295>`	Clears the particular BGP peer with the specified Autonomous System Number (ASN). The BGP peer ASN ranges from 1 to 4294967295.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Default

None.

Command Mode

Privileged Exec mode.

Applicability

This command was introduced in OcNOS version 6.3.4.

Examples

```
OcNOS#clear bgp 100 l2vpn evpn soft
OcNOS#clear bgp 100 l2vpn evpn soft in
```

clear bgp A.B.C.D l2vpn evpn

Use this command to reset the BGP session for neighbor with IPv4 address for L2VPN EVPN.

Command Syntax

```
clear bgp A.B.C.D l2vpn evpn soft (in | out | )
```

Parameters

`bgp A.B.C.D`	Clears the particular BGP neighbor with the specified IPv4 address.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Default

None.

Command Mode

Privileged Exec mode.

Applicability

This command was introduced in OcNOS version 6.3.4.

Examples

```
OcNOS#clear bgp 3.3.3.3 l2vpn evpn soft
OcNOS#clear bgp 3.3.3.3 l2vpn evpn soft in
```

clear bgp dampening

Use this command to reset BGP route flap dampening information.

Command Syntax

```
clear bgp ipv4 (unicast|multicast) dampening
clear bgp ipv4 (unicast|multicast) dampening A.B.C.D/M
clear ip bgp dampening
clear ip bgp dampening A.B.C.D/M
clear ip bgp ipv4 (unicast|multicast) dampening
clear ip bgp ipv4 (unicast|multicast) dampening A.B.C.D/M
```

Parameters

`ipv4`	IPv4 address family.
`multicast`	Multicast prefixes
`unicast`	Unicast prefixes
`A.B.C.D/M`	IP prefix (network/length), for example, 35.0.0.0/8

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp dampening 10.10.0.121/24
#clear ip bgp ipv4 unicast dampening
```

clear bgp external

Use this command to reset the BGP connection for all external peers.

Command Syntax

```
clear bgp external
clear bgp external in
clear bgp external in prefix-filter
clear bgp external out
clear bgp external soft
clear bgp external soft in
clear bgp external soft out
clear ip bgp external
clear ip bgp external in
clear ip bgp external in prefix-filter
clear ip bgp external out
clear ip bgp external soft
clear ip bgp external soft in
clear ip bgp external soft out
clear ip bgp external ipv4 (unicast|multicast) in
clear ip bgp external ipv4 (unicast|multicast) in prefix-filter
clear ip bgp external ipv4 (unicast|multicast) out
clear ip bgp external ipv4 (unicast|multicast) soft
clear ip bgp external ipv4 (unicast|multicast) soft in
clear ip bgp external ipv4 (unicast|multicast) soft out
```

Parameters

`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.
`ipv4`	Clear incoming advertised routes.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
`in`	Clear incoming advertised routes.
`prefix-filter`	

	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip bgp external
```

clear bgp flap-statistics

Use this command to reset BGP flap statistics.

Command Syntax

```
clear bgp ipv4 (unicast|multicast) flap-statistics
clear bgp ipv4 (unicast|multicast) flap-statistics A.B.C.D/M vrf (all | default |
  VRFNAME )
clear bgp ipv4 (unicast|multicast) flap-statistics A.B.C.D/M
clear ip bgp flap-statistics
clear ip bgp flap-statistics A.B.C.D/M
clear ip bgp ipv4 (unicast|multicast) flap-statistics
clear ip bgp ipv4 (unicast|multicast) flap-statistics A.B.C.D/M vrf (all | default
  | VRFNAME )
clear ip bgp ipv4 (unicast|multicast) flap-statistics A.B.C.D/M
```

Parameters

`ipv4`	IPv4 address family.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
`A.B.C.D/M`	IP prefix (network/length), for example, 35.0.0.0/8
`VRFNAME`	VPN routing or forwarding instance name
`all`	All VRF's
`default`	Default VRF

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp flap-statistics
#clear ip bgp ipv4 unicast flap-statistics
```

clear bgp peer-group

Use this command to reset the BGP connection for all members of a peer group.

Command Syntax

```
clear bgp peer-group WORD
clear bgp peer-group WORD in
clear bgp peer-group WORD in prefix-filter
clear bgp peer-group WORD out
clear bgp peer-group WORD soft
clear bgp peer-group WORD soft in
clear bgp peer-group WORD soft out
clear ip bgp peer-group WORD
clear ip bgp peer-group WORD in
clear ip bgp peer-group WORD in prefix-filter
clear ip bgp peer-group WORD out
clear ip bgp peer-group WORD soft
clear ip bgp peer-group WORD soft in
clear ip bgp peer-group WORD soft out
clear ip bgp peer-group WORD ipv4 (unicast|multicast) in
clear ip bgp peer-group WORD ipv4 (unicast|multicast) in prefix-filter
clear ip bgp peer-group WORD ipv4 (unicast|multicast) out
clear ip bgp peer-group WORD ipv4 (unicast|multicast) soft
clear ip bgp peer-group WORD ipv4 (unicast|multicast) soft in
clear ip bgp peer-group WORD ipv4 (unicast|multicast) soft out
```

Parameters

`WORD`	BGP peer-group name.
`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.
`ipv4`	Clear incoming advertised routes.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
	Clear incoming advertised routes.

`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp peer-group P1
```

clear bgp peer-group WORD l2vpn evpn

Use this command to reset the BGP L2VPN EVPN connection for all members of a BGP peer group.

Command Syntax

```
clear bgp peer-group WORD l2vpn evpn soft (in | out | )
```

Parameters

`bgp peer-group WORD`	Clears all the members of the specified peer group.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Default

None.

Command Mode

Privileged Exec mode.

Applicability

This command was introduced before OcNOS version 6.3.4.

Examples

```
OcNOS#clear bgp peer-group GRP1 l2vpn evpn soft
OcNOS#clear bgp peer-group GRP1 l2vpn evpn soft in
```

clear bgp statistics

Use this command to reset all BGP statistics.

Command Syntax

```
clear bgp statistics
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear bgp statistics
```

clear ip bgp * (description LINE|)*

Use this command to reset the BGP connection for all peers.

Command Syntax

```
clear ip bgp * (description LINE|)
```

Parameters

`description`	Administratively reset communication msg to neighbor

Command Mode

Privileged Exec mode

Applicability

This command is introduced in OcNOS version 6.0.0.

Examples

```
#clear bgp *
#clear bgp * description peer reset done for refresh
```

clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|)

Use this command to reset the BGP connection for peer.

Command Syntax

```
clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|)
```

Parameters

A.B.C.D	Address of the BGP neighbor in an IPv4 format
X:X::X:X	Address of the BGP neighbor in an IPv6 format
WORD	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
description	Administratively reset communication msg to neighbor

Command Mode

Privileged Exec mode

Applicability

This command is introduced in OcNOS version 6.0.0.

Examples

```
#clear ip bgp 40.1.1.2
#clear ip bgp 40.1.1.2 description peer reset is done
```

clear ip bgp A.B.C.D

Use this command to reset an IPv4 BGP neighbor address.

Command Syntax

```
clear ip bgp (A.B.C.D|WORD) in
clear ip bgp A.B.C.D in prefix-filter
clear ip bgp (A.B.C.D|WORD) out
clear ip bgp (A.B.C.D|WORD) soft
clear ip bgp (A.B.C.D|WORD) soft in
clear ip bgp (A.B.C.D|WORD) soft out
clear ip bgp A.B.C.D ipv4 (unicast|multicast) in
clear ip bgp A.B.C.D ipv4 (unicast|multicast) in prefix-filter
clear ip bgp A.B.C.D ipv4 (unicast|multicast) out
clear ip bgp A.B.C.D ipv4 (unicast|multicast) soft
clear ip bgp A.B.C.D ipv4 (unicast|multicast) soft in
clear ip bgp A.B.C.D ipv4 (unicast|multicast) soft out
```

Parameters

`A.B.C.D`	IPv4 address
`WORD`	Interface name
`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.
`ipv4`	Clear incoming advertised routes.
`multicast`	Multicast prefixes.
`unicast`	Unicast prefixes.
`in`	Clear incoming advertised routes.
`prefix-filter`	Push out prefix-list ORF and do inbound soft reconfig.
`out`	Clear outgoing advertised routes.
`soft`	Clear both incoming and outgoing routes.
`in`	Soft reconfig inbound update.
`out`	Soft reconfig outbound update.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp 35.0.0.1 in
```

clear ip bgp A.B.C.D vrf

Use this command to reset the VPN Routing/Forwarding (VRF) instance for a peer address.

Command Syntax

```
clear ip bgp A.B.C.D vrf WORD
clear ip bgp A.B.C.D vrf WORD in
clear ip bgp A.B.C.D vrf WORD out
clear ip bgp A.B.C.D vrf WORD soft
clear ip bgp A.B.C.D vrf WORD soft in
clear ip bgp A.B.C.D vrf WORD soft out
```

Parameters

`A.B.C.D`	IPv4 address
`WORD`	VPN routing/forwarding instance name
`in`	Clear incoming advertised routes
`out`	Clear outgoing advertised routes
`soft`	Clear both incoming and outgoing routes
`in`	Soft reconfig inbound update
`out`	Soft reconfig outbound update

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp 35.0.0.1 vrf
```

clear ip bgp all vrf (VRFNAME|all|default) (description LINE|)

Use this command to reset the BGP connection for vrf peer.

Command Syntax

```
clear ip bgp all vrf (VRFNAME|all|default) (description LINE|)
```

Parameters

`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF
`description`	Administratively reset communication msg to neighbor

Command Mode

Privileged Exec mode

Applicability

This command is introduced in OcNOS version 6.0.0

Examples

```
#clear ip bgp all vrf vrf1
#clear ip bgp all vrf all description all peer refresh is done
```

clear ip bgp table-map

Use this command to apply the modified table map or route map rules to the BGP routes in the existing IP routing table.

Command Syntax

```
clear ip bgp table-map (vrf (VRFNAME|all|default))
clear ip bgp ipv4 (unicast | multicast) table-map(vrf (VRFNAME|all|default))
```

Parameters

`vrf`	Select a VPN Routing/Forwarding Instance.
`VRFNAME`	Specify a VPN Routing/Forwarding instance name.
`all`	Select all VRFs.
`default`	Select default VRFs.
`unicast`	Unicast prefixes.
`multicast`	Multicast prefixes.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp table-map vrf all
```

clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|)

Use this command to reset the BGP connection for vrf peer.

Command Syntax

```
clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|)
```

Parameters

WORD	VPN routing/forwarding instance name
A.B.C.D	Address of the BGP neighbor in an IPv4 format
X:X::X:X	Address of the BGP neighbor in an IPv6 format
WORD	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
description	Administratively reset communication msg to neighbor

Command Mode

Privileged Exec mode

Applicability

This command is introduced in OcNOS version 6.0.0

Examples

```
#clear ip bgp vrf vrf1 10.1.1.2
#clear ip bgp vrf vrf1 10.1.1.2 description peer reset is done
```

debug bgp

Use this command to enable all BGP troubleshooting functions. Use this command without any parameters to turn on normal bgp debug information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug bgp (all|)
debug bgp bfd
debug bgp dampening
debug bgp events
debug bgp filters
debug bgp fsm
debug bgp keepalives
debug bgp nht
debug bgp nsm
debug bgp updates
debug bgp updates (in|out)
no debug bgp (all|)
no debug bgp bfd
no debug bgp dampening
no debug bgp events
no debug bgp filters
no debug bgp fsm
no debug bgp keepalives
no debug bgp nht
no debug bgp nsm
no debug bgp updates
```

Parameters

`all`	Used only with the `no` form; turns off all debugging for BGP
`bfd`	Enable debugging for BGP Bidirectional Forwarding Detection
`dampening`	Enable debugging for BGP dampening
`events`	Enable debugging for BGP events
`filters`	Enable debugging for BGP filters
`fsm`	Enable debugging for BGP Finite State Machine (FSM)
`keepalives`	Enable debugging for BGP keepalives
`nht`	Enable debugging for BGP NHT

`nsm`	Enable debugging for NSM messages
`updates`	Enable debugging for BGP updates
`in`	Debug inbound updates
`out`	Debug outbound updates

Command Mode

Privileged Exec mode and Configure Mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug bgp
#debug bgp events
```

distance bgp

Use this command to set BGP administrative distances. The administrative distance rates the trustworthiness of a routing information source. A higher distance value means a lower trust rating. For example, an administrative distance of 255 means that the routing information source cannot be trusted and should be ignored.

Use the no form of this command to restore the default value (110).

Command Syntax

```
distance <1-255>
distance <1-255> A.B.C.D/M (WORD|)
no distance <1-255>
no distance <1-255> A.B.C.D/M (WORD|)
```

Parameters

`<1-255>`	Distance for BGP routes
`A.B.C.D/M`	Distance for routes to prefixes whose nexthop matches this address.
`WORD`	Name of access list to apply to route updates.

Command Mode

Router BGP mode

Defaults

None

Applicability

This command was introduced in OcNOS version 6.0.0.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#distance 100 10.10.10.10/30
```

exit-address-family

Use this command to exit Address-Family-vrf, Address-Family 7mode.

For information on how to enter the address family mode (IPv4), see address-family.

Command Syntax

```
exit-address-family
```

Parameters

None

Default

No default value is specified

Command Mode

Address Family-vrf

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following examples shows the change in the prompt after using this command.

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family ipv4 multicast
(config-router-af)#exit-address-family
(config-router)#
```

ip as-path access-list

Use this command to define a BGP Autonomous System (AS) path access list. A named community list is a filter based on regular expressions. If the regular expression matches the specified string representing the AS path of the route, then the permit or deny condition applies. Use this command to define the BGP access list globally; use the neighbor router configuration command to apply a specific access list.

Use the no parameter with this command to disable use of the access list.

Command Syntax

```
ip as-path access-list WORD (deny|permit) LINE
no ip as-path access-list WORD
no ip as-path access-list WORD (deny|permit) LINE
```

Parameters

`WORD`	Access list name
`deny`	Reject packets
`permit`	Forward packets
`LINE`	An ordered list as a regular expression

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip as-path access-list mylist deny ^65535$
```

ip community-list <1-99>

Use this command to specify a standard community list (1 to 99) that specifies BGP community attributes.

Use the `no` parameter with this command to delete the community list entry.

Command Syntax

```
ip community-list <1-99> (deny|permit)
ip community-list <1-99> (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
no ip community-list <1-99> (deny|permit)
no ip community-list <1-99> (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
```

Parameters

`deny`	Reject the community
`permit`	Accept the community
`AA:NN`	Community number
`internet`	Advertise routes to the internet community
`local-AS`	Do not advertise routes to external BGP peers
`no-advertise`	Do not advertise routes to other BGP peers
`no-export`	Do not advertise routes outside of Autonomous System boundary

Default

By default, ip community list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip community-list 55 permit 7675:80 7675:90

(config)#no ip community-list 55 permit 7675:80 7675:90
```

ip community-list <100-500>

Use this command to specify an expanded community list (100 to 500) that specifies BGP community attributes.

Use the `no` parameter with this command to delete the community list entry.

Command Syntax

```
ip community-list <100-500> (deny|permit)
ip community-list <100-500> (deny|permit) LINE
no ip community-list <100-500>
no ip community-list <100-500> (deny|permit) LINE
```

Parameters

`deny`	Reject community
`permit`	Accept community
`LINE`	An ordered list as a regular expression

Default

By default, ip community list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip community-list 225 permit 6789906
(config)#ip community-list expanded CLIST permit .*
```

ip community-list expanded

Use the community-lists to specify BGP community attributes. The community attribute is used for implementing policy routing. It is an optional, transitive attribute and facilitates transfer of local policies through different autonomous systems. It includes community values that are 32-bits long.

There are two kinds of community-lists: expanded and standard. The standard community-list defines the community attributes in a specified format and not with regular expressions. The expanded community-list defines the community attributes with regular expressions.Use the `no` parameter with this command to delete the community list entry.

Command Syntax

```
ip community-list expanded WORD (deny|permit) LINE
no ip community-list expanded WORD
no ip community-list expanded WORD (deny|permit) LINE
```

Parameters

`WORD`	Community list name
`deny`	Reject community
`permit`	Accept community
`LINE`	An ordered list as a regular expression

Default

By default, ip community list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip community-list 125 permit 6789906
(config)#ip community-list expanded CLIST permit .*
```

ip community-list standard

Use the community-lists to specify BGP community attributes. The community attribute is used for implementing policy routing. It is an optional, transitive attribute and facilitates transfer of local policies through different autonomous systems. It includes community values that are 32-bits long. There are two kinds of community-lists: expanded and standard. The standard community-list defines the community attributes in a specified format without regular expressions. The expanded community-list defines the community attributes with regular expressions.

Use this command to add a standard community-list entry. The standard community-list is compiled into binary format and is directly compared with the BGP communities attribute in the BGP updates. The comparison is faster than the expanded community-list. Any community value that does not match the standard community value is automatically treated as expanded.

Use the `no` parameter with this command to delete the standard community-list entry.

Command Syntax

```
ip community-list standard WORD (deny|permit)
ip community-list standard WORD (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
no ip community-list standard WORD
no ip community-list standard WORD (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
```

Parameters

`WORD`	Community list name
`deny`	Rejects a specified community when a community value is provided. If no community value is configured, it will reject all prefixes or routes listed in the community-list.
`permit`	Accepts a specified community when a community value is provided. If no community value is configured, it will accept all prefixes or routes listed in the community-list.
`AA:NN`	Community number
`internet`	Advertise routes to the internet community
`local-AS`	Do not advertise routes to external BGP peers
`no-advertise`	Do not advertise routes to other BGP peers
`no-export`	Do not advertise routes outside of Autonomous System boundary

Default

By default, ip community list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip community-list standard CLIST permit 7675:80 7675:90 no-export
```

ip community-list WORD

Use the community-list commands to specify BGP community attributes. The community attribute is used for implementing policy routing. It is an optional, transitive attribute and facilitates transfer of local policies through different autonomous systems. There are two kinds of community-lists: the expanded and standard. The `standard community-list` defines the community attributes in a specified format and not with regular expressions. The `expanded community-list` defines the community attributes with regular expressions.

Use the `no` parameter with this command to delete the community list entry.

Command Syntax

```
ip community-list WORD (deny|permit)
ip community-list WORD (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
no ip community-list WORD
no ip community-list WORD (deny|permit) [AA:NN|internet|local-AS|no-advertise|no-export]
```

Parameters

`WORD`	Community list name
`deny`	Rejects a specified community when a community value is provided. If no community value is configured, it will reject all prefixes or routes listed in the community-list.
`permit`	Accepts a specified community when a community value is provided. If no community value is configured, it will accept all prefixes or routes listed in the community-list.
`AA:NN`	Community number
`internet`	Advertise routes to the internet community
`local-AS`	Do not advertise routes to external BGP peers
`no-advertise`	Do not advertise routes to other BGP peers
`no-export`	Do not advertise routes outside of Autonomous System boundary

Default

By default, ip community list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip community-list mylist permit 7675:80 7675:90

(config)#no ip community-list mylist permit 7675:80 7675:90
```

ip extcommunity-list <1-99>

Use this command to create an entry for a standard extended community list.

Use the `no` parameter with this command to delete the community-list entry.

Command Syntax

```
ip extcommunity-list <1-99> (deny|permit) LINE (rt|soo)
no ip extcommunity-list <1-99> (deny|permit) LINE (rt|soo)
```

Parameters

`deny`	Reject community
`permit`	Accept community
`LINE`	One of the following:
`rt`	Route target extended community in aa:nn or IPaddr:nn format
`soo`	Site-of-origin extended community in aa:nn or IPaddr:nn format

Default

By default, ip extcommunity list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip extcommunity-list 3 permit rt 10.10.23.123:67

(config)#ip extcommunity-list 25 deny soo 1465:22
```

ip extcommunity-list <100-500>

Use this command to create an extended community list.

Use the `no` parameter with this command to delete the community-list entry.

Command Syntax

```
ip extcommunity-list <100-500> (deny|permit) LINE
no ip extcommunity-list <100-500> (deny|permit) LINE
```

Parameters

`<100-500>`	Extended community list number (expanded)
`deny`	Reject the community
`permit`	Accept the community
`LINE`	Any regular expression:

Default

By default, ip extcommunity list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip extcommunity-list 125 permit 4567:335

(config)#ip extcommunity-list 231 deny *.
```

ip extcommunity-list expanded

Use this command to create an expanded community list.

Use the `no` parameter with this command to delete the expanded community-list entry.

Command Syntax

```
ip extcommunity-list expanded WORD
ip extcommunity-list expanded WORD (deny|permit) LINE
no ip extcommunity-list expanded WORD
no ip extcommunity-list expanded WORD (deny|permit) LINE
```

Parameters

`WORD`	Expanded community list name
`deny`	Reject the community
`permit`	Accept the community
`LINE`	One of the following:
`rt`	Route target extended community in aa:nn or IPaddr:nn format
`soo`	Site-of-origin extended community in aa:nn or IPaddr:nn format

Default

By default, ip extcommunity list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip extcommunity-list 125 permit 4567335
(config)#ip extcommunity-list expanded CLIST permit .*
```

ip extcommunity-list standard

Use this command to create and delete a standard extended-community list. The extended community attribute is 8 bytes in 2 formats. The sub-type can be route target (`rt`) or site of origin (`soo`). Thus, the sub-type of each community must be specified when creating the extended community list. Regarding the formats, an extended community is based on a 6-byte value. These 6-bytes are represented in 4-btye:2-byte format, and may be entered in one of the following forms:

- Format 1, aa.nn: The 16-bit value of the AS (aa) number is represented in the higher-order 4-bytes. If the extended ASN capability is enabled, the AS number is represented using higher-order 4-bytes. The nn assigned value is represented in the low-order 2-bytes in both cases.
- Format 2, IPaddr:nn: In this format, the higher-order 4-bytes are used to represent the IP address, and the low-order 2-bytes are used to represent the assigned value.

Use the `no` parameter with this command to delete the extended-community-list entry.

Command Syntax

```
ip extcommunity-list standard WORD (deny|permit) (rt|soo) (aa:nn)
no ip extcommunity-list standard WORD (deny|permit) (rt|soo) (aa:nn)
```

Parameters

`WORD`	Extended community list name
`deny`	Reject the community
`permit`	Accept the community
`rt`	Route target extended community in aa:nn or IPaddr:nn format
`soo`	Site-of-origin extended community in aa:nn or IPaddr:nn format

Default

By default, ip extcommunity list is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip extcommunity-list standard 36 permit rt 5675:50
(config)#ip extcommunity-list standard CLIST permit soo 10.10.32.15:70
```

l2vpn-unnumbered-mode

Use this command to configure the unnumbered peer with all the neighbor configs applicable in L2VPN EVPN address family.

Use the `exit` parameter with this command to leave L2VPN unnumbered mode.

Command Syntax

```
bgp l2vpn-unnumbered-mode
exit-l2vpn-unnumbered-mode
```

Parameters

None

Applicability

No default value is specified

Command Mode

Address-family IPv4 Unicast Mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-l2vpn-unnum)#
```

match ip peer

Use this command to apply policies based on the route source of which the BGP TCP/IP session formed using an IPv4 address in the update message.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
match ip peer (<1-199>|<1300-2699>|WORD)
no match ip peer (<1-199>|<1300-2699>|WORD)
```

Parameters

`<1-199>`	IP access-list number
`<1300-2699>`	IP access-list number (expanded range)
`WORD`	Access-list name

Default

By default, import bgp route is disabled

Command Mode

Route-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#route-map in-A permit 10
(route-map)#match ip peer 1
```

match large-community

Use this command to match the large community.

Command Syntax

```
match large-community XXX:YYYY:ZZZ
no match large-community XXX:YYYY:ZZZ
```

Parameters

XXX:YYYY:ZZZ	Large community number

Default

By default send-community large is enabled for peer

Command Mode

Route-map mode

Applicability

This command was introduced in OcNOS version 6.1.0.

Examples

```
(config)#route-map R1 permit 10
(config-route-map)#match large-community 1111:2222:33333
(config-route-map)#
```

max-paths

Use this command to set the number of equal-cost multi-path (ECMP) routes for eBGP or iBGP. You can install multiple BGP paths to the same destination to balance the load on the forwarding path.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
max-paths (ebgp|ibgp|) <2-64>
no max-paths ebgp (<2-64>|)
no max-paths ibgp (<2-64>|)
```

Parameters

`ebgp`	eBGP ECMP session
`ibgp`	iBGP ECMP session
`<2-64>`	Number of routes

Default

Available for the IPv4 and IPv6 unicast addresses.

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example configures 7 routes for ECMP for iBGP.

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv4 unicast
(config-router-af)#max-paths ibgp 7
```

max-paths eigbp

Use this command to set the number of equal-cost multi-path (ECMP) routes for both eBGP and iBGP. This feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in BGP networks.

Use the `no` parameter with this command to disable this feature and set the BGP path selection to the default single best path.

Command Syntax

```
max-paths eibgp <2-64>
no max-paths eibgp (<2-64>|)
```

Parameters

`<2-64>`	Number of routes

Default

Not applicable

Command Mode

IPv4 and IPv6 unicast address-family mode

Applicability

This command was introduced in OcNOS version 5.0.

Exception with Administrative Distance

Administrative distance (AD) is the feature that routers use to select the best path when there are two or more different routes to the same destination from two different routing protocols.

In OcNOS, AD is per route. With max-paths eibgp, BGP programs the AD(200)/AD20 for eiBGP routes. If there is a route from IGP with a lower AD, it will be preferred and get installed in the FIB.

A BGP route cannot have a different AD for each path. Multipath candidates of eibgp max-paths can be programmed either with AD20 or AD200.

Currently AD20/AD200 will be programmed for all eiBGP routes as per best path selected by BGP. For example, if the best path is learned via iBGP peering, then AD200 is programmed for all the installed multipaths (including paths from eBGP peering) or If the best path is learned via eBGP peering, then AD20 is programmed for all the installed multipaths (including paths from iBGP peering).

Example

This example configures the number of parallel iBGP and eBGP routes that can be installed into a routing table.

```
(config)#router bgp 100
(config-router)#neighbor 11.0.0.2 remote-as 100
(config-router)#neighbor 12.0.0.1 remote-as 200
(config-router)#neighbor 1101::2 remote-as 200
(config-router)#neighbor 1201::1 remote-as 100
(config-router)#address-family ipv4 unicast
(config-router-af)#max-paths eibgp 2
```

```
(config-router-af)#network 11.0.0.2 activate
(config-router-af)#network 12.0.0.1 activate
(config-router-af)#exit-address-family
(config-router)#address-family ipv6 unicast
(config-router-af)#max-paths eibgp 3
(config-router-af)#network 1101::2 activate
(config-router-af)#network 1201::1 activate
(config-router-af)#exit-address-family
```

neighbor activate

Use this command to enable the exchange of specific address family routes with a neighboring router. After a TCP connection is opened with a neighboring router, use this command to enable or disable the exchange of address family information. To enable the exchange of multicast, activate neighbors using this command in address family mode.

Use the `no` parameter with this command to disable exchange of information with a neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) activate
no neighbor (A.B.C.D|X:X::X:X|WORD) activate
```

For L2VPN-unnumbered mode:

```
neighbor WORD activate
no neighbor WORD activate
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default activate is disabled. For all address-families, use this command to enable a neighbor to exchange routing information of a specific address-family with a neighbor.

Command Mode

Address Family mode and l2vpn-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router-af)#neighbor 10.10.20.1 activate

(config)#router bgp 100
(config-router)#address-family l2vpn evpn
(config-router-af)#bgp l2vpn-unnumbered-mode
(config-router-l2vpn-unnum)#neighbor eth1 activate
```

neighbor advertisement-interval

Use this command to set a minimum interval between the sending of BGP routing updates. To reduce the flapping of routes, set a minimum advertisement interval so that the BGP routing updates are sent only per interval seconds. BGP dampening can also be used to control the effects of flapping routes.

Use the `no` parameter with this command to set the interval time to default.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) advertisement-interval <0-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) advertisement-interval
no neighbor (A.B.C.D|X:X::X:X|WORD) advertisement-interval <0-65535>
```

For BGP ununbered mode:

```
neighbor WORD advertisement-interval <0-65535>
no neighbor WORD advertisement-interval
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<0-65535>`	Advertisement interval value in seconds

Default

By default, neighbor value for ebgp peer is 30 seconds and IBGP peer is 5 seconds

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.3 advertisement-interval 45
```

For unnumbered peer this configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 advertisement-interval 20
```

neighbor allowas-in

Use this command to advertise prefixes (routes) even when the source of the prefixes is from the same Autonomous System (AS) number.

Use this command in a scenario where two routers at different locations use the same Autonomous System number and are connected via an ISP. Once prefixes arrive from one branch at the ISP, they are tagged with the customer's AS number. By default, when the ISP passes the prefixes to the other router, the prefixes are dropped if the other router uses the same AS number. Use this command to advertise the prefixes at the other side. Control the number of times an AS number is advertised by specifying a number.

In a hub and spoke configuration in a VPN, a PE (Provider Edge) router advertises all prefixes containing duplicate AS numbers. Use this command to configure two VRFs on each PE router to receive and advertise prefixes. One of the VRFs receives prefixes with AS numbers from all PE routers and then advertises them to neighboring PE routers. The other VRF receives prefixes with AS numbers from the CE (Customer Edge) router and advertises them to all PE routers in the hub and spoke configuration.

Use the `no` parameter with this command to remove the configuration.

For unnumbered peer IPv4 unicast mode commands are configured under v4-unnumbered-mode.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) allowas-in <1-10>
no neighbor (A.B.C.D|X:X::X:X|WORD) allowas-in
```

For v4-unnumbered mode:

```
neighbor WORD allowas-in
neighbor WORD allowas-in <1-10>
no neighbor WORD allowas-in
```

Parameters

`A.B.C.D`	IPv4 neighbor address.
`X:X::X:X`	IPv6 neighbor address.
`WORD`	Name of peer group.

Note: For information on how to create peer groups, refer to the neighbor peer-group and neighbor remote-as commands. When this parameter is used with a command, the command applies on all peers in the specified group.

`<1-10>`	Number of times to allow the advertisement of an AS number

Default

No default value is specified.

Command Mode

This CLI command is available to configure in all Address-family modes (v4-unnumbered mode, IPv4 unicast, IPv6 unicast, IPv4 labeled-unicast, IPv6 labeled-unicast, VPNv6 unicast, IPv4/IPv6 VRF, L2VPN EVPN)

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 10.10.0.1 allowas-in 3

#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 allowas-in 3

#con terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 allowas-in 6
```

For unnumbered peer configs in L2vpn evpn mode are configured under l2vpn-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor eth1 allowas-in 7
```

neighbor as-origination-interval

Use this command to set the minimum interval between sending of AS-origination routing updates.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) as-origination-interval <1-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) as-origination-interval
no neighbor (A.B.C.D|X:X::X:X|WORD) as-origination-interval <1-65535>
```

For BGP unnumbered mode:

```
neighbor WORD as-origination-interval interval <0-65535>
no neighbor WORD as-origination-interval
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-65535>`	AS origination interval in seconds

Default

By default, neighbor as origination interval is 15 seconds

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.75 as-origination-interval 555
```

For unnumbered peer this configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 as-origination-interval 50
```

neighbor attribute-unchanged

Use this command to advertise unchanged BGP attributes to the specified neighbor.

Use the `no` parameter with this command to disable this function.

For unnumbered peer IPv4 unicast mode commands are configured under v4-unnumbered-mode.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) attribute-unchanged ({ as-path|next-hop|med }|)
no neighbor (A.B.C.D|X:X::X:X|WORD) attribute-unchanged (({ as-path|next-hop|
  med }|)
```

For v4-unnumbered mode:

```
neighbor WORD attribute-unchanged ({ as-path|next-hop|med }|)
no neighbor WORD attribute-unchanged (({ as-path|next-hop|med }|)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`as-path`	AS path attribute
`next-hop`	Nexthop attribute
`med`	Multi-exit discriminator attribute

Default

By default, the neighbor attribute-unchanged is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.75 attribute-unchanged as-path med

(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#OcNOS(config-router-v4-unnum)#neighbor eth1
attribute-unchanged as-path
```

neighbor authentication-key

Use this command to enable message digest5 (MD5) authentication on a TCP connection between two BGP peers. Configuring MD5 authentication between two BGP peers, means that each segment sent on the TCP connection between the peers is verified. MD5 authentication must be configured with the same password on both BGP peers; otherwise, the connection between them will not be established.

Use the `no` parameter with this command to delete the MD5 authentication.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) authentication-key (0 | 1 |) <WORD>
<WORD> plain text password
neighbor (A.B.C.D|X:X::X:X|WORD) authentication-key <WORD>
<WORD> encrypted password
no neighbor (A.B.C.D|X:X::X:X|WORD) authentication-key <WORD>
<WORD> encrypted password
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of the BGP peer group
`0`	Unencrypted password
`WORD`	Unencrypted password (maximum length 80 characters)
`1`	Encrypted password (default)
`WORD`	Encrypted password (maximum length 162 characters)

Default

Not applicable

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 11
(config-router)#neighbor 10.10.0.73 authentication-key 0 myPass
(config-router)#no neighbor 10.10.0.73 authentication-key 0xba76ef302e2f20af
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 authentication-key 0 key1
```

neighbor capability graceful-restart

Use this command to advertise the graceful restart capability to its neighbor. This configuration indicates that the BGP speaker has the ability to preserve its forwarding state for the address family when BGP restarts.

You must first specify a neighbor's `remote-as` identification number assigned by the neighbor router.

Note: The graceful restart capability is advertised only when the graceful restart capability has been enabled using the bgp graceful-restart command.

Use the `no` parameter with this command to not advertise the graceful restart capability to its neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) capability graceful-restart
no neighbor (A.B.C.D|X:X::X:X|WORD) capability graceful-restart
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, the graceful-restart is disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.10.50 capability graceful-restart
```

neighbor capability orf prefix-list

Use this command to enable Outbound Router Filtering (ORF) and advertise the ORF capability to its neighbors. The ORFs send and receive capabilities to lessen the number of updates exchanged between neighbors. By filtering updates, this option minimizes generating and processing of updates.

The two routers exchange updates to maintain the ORF for each router:

- The local router advertises the ORF capability in `send` mode.
- The remote router receives the ORF capability in `receive` mode, applying the filter as outbound policy.

Only an individual router or a peer group can be configured to be in `receive` or `send` mode. A peer-group member cannot be configured to be in `receive` or `send` mode.

Use the `no` parameter with this command to disable this function.

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) capability orf prefix-list (both|receive|send)
no neighbor (A.B.C.D|X:X::X:X|WORD) capability orf prefix-list (both|receive|send)
```

For v4-unnumbered mode:

```
neighbor WORD capability orf prefix-list (both|receive|send)
no neighbor WORD capability orf prefix-list (both|receive|send)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`both`	The local router can send ORF entries to its peer, as well as receive ORF entries from its peer.
`receive`	The local router is willing to receive ORF entries from its peer
`send`	The local router is willing to send ORF entries to its peer

Default

By default, the orf prefix-list is disabled

Command Mode

Router Address Family (IPv4 unicast, IPv4 multicast, IPv6, v4-unnumbered) mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
```

```
(config-router-af)#neighbor 10.10.0.5 capability orf prefix-list both
(config-router-af)#neighbor effe:2897::0003:3ed5 capability orf prefix-list
receive

(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 capability orf prefix-list both
```

neighbor collide-established

Use this command to include a neighbor already in an established state for conflict resolution when a TCP connection collision is detected. This command is not required for most network deployments, so users should only use this command when required.

Note: The associated functionality of including an “established” neighbor into TCP connection collision conflict resolution is automatically enabled when a neighbor is configured for BGP graceful-restart.

Use the `no` option with this command to turn this feature off.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) collide-established
no neighbor (A.B.C.D|X:X::X:X|WORD) collide-established
```

For BGP unnumbered mode:

```
neighbor WORD collide-established
no neighbor WORD collide-established
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, neighbor collide is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 3.3.3.3 collide-established
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 collide-established
```

neighbor de-activate

Use this command to disable the exchange of specific address family routes with a neighboring router for peer-group member. After a TCP connection is opened with a neighboring router, use this command to enable or disable the exchange of address family information for peer-group member.

Use the `no` parameter with this command to inherit the properties of peer-group. i.e., If peer-group is activated in the given address-family, no of this command will activate the peer member as peer-group is activated in address-family..

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X) de-activate
no neighbor (A.B.C.D|X:X::X:X) de-activate
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format

Default

By default de-activate is disabled **for all peer-group members.** For all address-families, Use this command to disable a neighbor to exchange routing information of a specific address-family with a neighbor.

Command Mode

Address Family mode

Applicability

This command is introduced in OcNOS version 6.0.0. This command is applicable only to neighbor which is a member of peer-group.

Examples

```
#configure terminal
router bgp 100
neighbor pg1 peer-group
neighbor pg1 remote-as 100
neighbor 10.1.1.2 peer-group pg1
neighbor 20.1.1.2 peer-group pg1
neighbor 30.1.1.2 remote-as 100
!
address-family ipv4 unicast
neighbor pg1 activate
neighbor 30.1.1.2 activate
exit-address-family
!
(config-router-af)#neighbor 10.1.1.2 de-activate
(config-router-af)#show running-config bgp
!
router bgp 100
neighbor pg1 peer-group
neighbor pg1 remote-as 100
neighbor 10.1.1.2 peer-group pg1
neighbor 20.1.1.2 peer-group pg1
```

```
neighbor 30.1.1.2 remote-as 100
!
address-family ipv4 unicast
neighbor pg1 activate
neighbor 10.1.1.2 de-activate
neighbor 30.1.1.2 activate
exit-address-family
!
(config-router-af)#no neighbor 10.1.1.2 de-activate
(config-router-af)#commit
(config-router-af)#
(config-router-af)#show running-config bgp
!
router bgp 100
neighbor pg1 peer-group
neighbor pg1 remote-as 100
neighbor 10.1.1.2 peer-group pg1
neighbor 20.1.1.2 peer-group pg1
neighbor 30.1.1.2 remote-as 100
!
address-family ipv4 unicast
neighbor pg1 activate
neighbor 30.1.1.2 activate
exit-address-family
!
```

neighbor default-originate

Use this command to allow a BGP local router to send the default route 0.0.0.0 to a neighbor to use as a default route. This command can be used with standard or extended access lists.

Use the `no` parameter with this command to send no route as a default.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) default-originate
neighbor (A.B.C.D|X:X::X:X|WORD) default-originate route-map WORD
no neighbor (A.B.C.D|X:X::X:X|WORD) default-originate
no neighbor (A.B.C.D|X:X::X:X|WORD) default-originate route-map WORD
```

For v4-unnumbered mode:

```
neighbor WORD default-originate
neighbor WORD default-originate route-map WORD
no neighbor WORD default-originate
no neighbor WORD default-originate route-map WORD
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Route map name

Default

By default, neighbor default originate is disabled

Command Mode

Route Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.10.1 default-originate route-map myroute
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
#con terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 default-originate
```

neighbor description

Use this command to associate a description with a neighbor. This command helps to identify a neighbor quickly. This command is useful for an ISP that has multiple neighbor relationships.

Use the `no` parameter with this command to remove the description.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) description LINE
no neighbor (A.B.C.D|X:X::X:X|WORD) description
no neighbor (A.B.C.D|X:X::X:X|WORD) description LINE
```

For BGP unnumbered mode:

```
neighbor WORD description LINE
no neighbor WORD description
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`LINE`	Neighbor description (up to 80 characters)

Default

By default, the neighbor description is disabled

Command Mode

Router mode and Address Family and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 1.2.3.4 description Backup router for sales

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 description Bank of America
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 description unnbered-peer1
```

neighbor disallow-infinite-holdtime

Use this command to disallow configuration of infinite hold-time. This command enables the local BGP speaker to reject a hold-time of “0” seconds from a peer (during exchange of open messages) or a user (during configuration).

The `no` form of this command allows the BGP speaker to accept a hold-time of “0” from a peer or during configuration.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) disallow-infinite-holdtime
no neighbor (A.B.C.D|X:X::X:X|WORD) disallow-infinite-holdtime
```

For BGP unnumbered mode:

```
neighbor WORD disallow-infinite-holdtime
no neighbor WORD disallow-infinite-holdtime
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, neighbor disallow infinite holdtime is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config-router)#neighbor 10.11.4.26 disallow-infinite-holdtime
(config-router)#neighbor 3ffe::45 disallow-infinite-holdtime
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 disallow-infinite-holdtime
```

© 2025 IP Infusion Inc. Proprietary

neighbor distribute-list

Use this command to filter route updates from a particular BGP neighbor. Use only one distribute list per BGP neighbor.

Use the `no` parameter with this command to remove an entry.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) distribute-list (<1-199>|<1300-2699>|WORD)
  (in|out)
no neighbor (A.B.C.D|X:X::X:X|WORD) distribute-list (<1-199>|<1300-2699>|WORD)
  (in|out)
```

For v4-unnumbered mode:

```
neighbor WORD distribute-list (<1-199>|<1300-2699>|WORD) (in|out)
no neighbor WORD distribute-list (<1-199>|<1300-2699>|WORD) (in|out)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-199>`	IP access-list number
`<1300-2699>`	IP access-list number (expanded-range)
`WORD`	Access-list name
`in`	Filter incoming advertised routes
`out`	Filter outgoing advertised routes

Default

By default, neighbor distribute list is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 1.2.3.4 distribute-list mylist out
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 distribute-list list1 in
```

© 2025 IP Infusion Inc. Proprietary

neighbor dont-capability-negotiate

Use this command to disable capability negotiation. This command is used to allow compatibility with older BGP versions that have no capability parameters used in open messages between peers.

Use the `no` parameter with this command to enable capability negotiation.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) dont-capability-negotiate
no neighbor (A.B.C.D|X:X::X:X|WORD) dont-capability-negotiate
```

For BGP unnumbered mode:

```
neighbor WORD dont-capability-negotiate
no neighbor WORD dont-capability-negotiate
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, capability negotiation is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.34 dont-capability-negotiate
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 dont-capability-negotiate
```

neighbor ebgp-multihop

Use this command to accept and try BGP connections to external peers on indirectly connected networks. Multihop is not established if the only route to the multihop peer is a default route. This avoids loop formation.

Use the `no` parameter with this command to return to the default.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) ebgp-multihop
neighbor (A.B.C.D|X:X::X:X|WORD) ebgp-multihop <1-255>
no neighbor (A.B.C.D|X:X::X:X|WORD) ebgp-multihop
no neighbor (A.B.C.D|X:X::X:X|WORD) ebgp-multihop <1-255>
```

For BGP unnumbered mode:

```
neighbor WORD ebgp-multihop
neighbor WORD ebgp-multihop <1-255>
no neighbor WORD ebgp-multihop
no neighbor WORD ebgp-multihop <1-255>
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-255>`	Maximum hop count

Default

By default, maximum hop count is 255

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.10.34 remote-as 20
(config-router)#neighbor 10.10.10.34 ebgp-multihop 5
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 ebgp-multihop
```

neighbor enforce-multihop

Use this command to enforce BGP neighbors to perform multihop.

Use the `no` parameter with this command to turn off this feature.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) enforce-multihop
no neighbor (A.B.C.D|X:X::X:X|WORD) enforce-multihop
```

For BGP unnumbered mode:

```
neighbor WORD enforce-multihop
no neighbor WORD enforce-multihop
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, the maximum hop count is 255

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.34 remote-as 20
(config-router)#neighbor 10.10.0.34 enforce-multihop
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 enforce-multihop
```

© 2025 IP Infusion Inc. Proprietary

neighbor extended-optional-param

Use this command enable extended optional parameter length for BGP OPEN messages as specified in RFC 9072.

Use the `no` form of this command to disable extended optional parameter length.

Command Syntax

```
neighbor A.B.C.D|X:X::X:X|WORD extended-optional-param
no neighbor A.B.C.D|X:X::X:X|WORD extended-optional-param
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

Extended optional parameter length is disabled by default.

Command Mode

BGP router mode

Applicability

This command was introduced in OcNOS version 6.0.0.

Example

```
(config)#router bgp 100
(config-router)#neighbor 1.1.1.1 extended-optional-param
(config-router)#commit
(config-router)#show run bgp
!
router bgp 100
bgp router-id 11.11.11.11
neighbor 1.1.1.1 remote-as 100
neighbor 1.1.1.1 extended-optional-param
```

neighbor fall-over bfd

Use this command to enable the BFD for BGP peers.

Use the `no` form of the command to disable the BFD option for BGP peers.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) fall-over bfd (multihop|)
no neighbor (A.B.C.D|X:X::X:X|WORD) fall-over bfd (multihop|)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format.
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`multihop`	Enable multihop

Default

Bidirectional Forwarding Detection for BGP peers is disabled by default.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
(config)#router bgp 100
(config-router)#neighbor 10.2.3.4 fall-over bfd
(config-router)#
(config-router)#neighbor 10.2.3.4 fall-over bfd multihop
```

neighbor filter-list

Use this command to set up a BGP filter. This command specifies an access list filter on updates based on the BGP autonomous system paths. Each filter is an access list based on regular expressions.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) filter-list WORD (in|out)
no neighbor (A.B.C.D|X:X::X:X|WORD) filter-list WORD (in|out)
```

For v4-unnumbered mode:

```
neighbor WORD filter-list WORD (in|out)
no neighbor WORD filter-list WORD (in|out)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Name of an autonomous system path access list
`in`	Filter incoming advertised routes
`out`	Filter outgoing advertised route

Default

By default, filter list is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.34 remote-as 20
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.34 filter-list out in
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 filter-list list2 in
```

neighbor g-shut

Use this command to start a graceful shutdown for the BGP session of the specified BGP neighbor. The BGP session for this neighbor is shut down after the graceful shutdown timer expires.

If there is no alternate path available for traffic to flow prior the actual shutdown of the BGP session, this path is made available for 60 seconds or for configured time after which the path is no longer available and traffic is dropped.

Use the `no` parameter with this command to bring up the session again for the specified BGP neighbor whose BGP session had been shut down using the `neighbor g-shut` command.

Note: The graceful shutdown capability is not supported on iBGP sessions.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) g-shut
no neighbor (A.B.C.D|X:X::X:X|WORD) g-shut
```

Parameters

`A.B.C.D`	Neighbor IPv4 address
`X:X::X:X|`	Neighbor IPv6 address
`WORD`	Neighbor tag

Default

By default, neighbor g-shut is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#neighbor 1.1.1.2 g-shut

#configure terminal
(config)#router bgp 100
(config-router)#no neighbor 1.1.1.2 g-shut
```

neighbor g-shut-timer

Use this command to configure the value of the graceful shutdown timer. After the timer expires, the BGP session initiated for graceful shutdown is shut down.

Use the `no` parameter with this command to revert to the default setting.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) g-shut-timer <10-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) g-shut-timer <10-65535>
```

Parameters

`A.B.C.D`	Neighbor IPv4 address
`X:X::X:X`	Neighbor IPv6 address
`WORD`	Neighbor tag
`<10-65535>`	Graceful shutdown timer in seconds

Default

By default, the timer value is set to 60 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#neighbor 1.1.1.2 g-shut-timer 120
```

neighbor limit

Use this command to specify the maximum number of peers that can be configured in the BGP dynamic peer-group.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor WORD limit <1-200>
no neighbor WORD (limit <1-200>|)
```

Parameters

`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command.
`<1-200>`	The maximum number of peers that can be configured in a BGP dynamic peer-group.

Default

By default, neighbor word limit is disabled

Command Mode

Router mode and Address Family VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 11
(config-router)#neighbor group1 limit 120
```

neighbor local-as

Use this command to specify an AS (autonomous system) number to use with BGP neighbor.

Use the `no` parameter with this command to disable this command.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295>  (no-prepend|) (replace-as|)
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295>
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295> no-prepend
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295> replace-as
```

For BGP unnumbered mode:

```
neighbor WORD local-as <1-4294967295>  (no-prepend|) (replace-as|)
no neighbor WORD local-as <1-4294967295>
no neighbor WORD local-as <1-4294967295> no-prepend
no neighbor WORD local-as <1-4294967295> replace-as
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-4294967295>`	Neighbor's AS number when extended capabilities are configured
`no-prepend`	Do not prepend local-as to updated from ebgp peers
`replace-as`	Replace real AS with local-as in the EBGP updates

Note: The AS number 23456 is a reserved 2-byte AS number. An old BGP speaker (2-byte implementation) should be configured with 23456 as its remote AS number while peering with a non-mappable new BGP speaker (4-byte implementation).

Default

By default, local-as is disabled.

Command Mode

Router mode and Address Family-vrf mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3. The new version of the command with “no-prepend” and “replace-as“ option

is introduced in OcNOS version 6.4.1.

Example

```
#configure terminal
(config)#router bgp 100
(config-router)#neighbor 20.1.1.3 remote-as 300
(config-router)#neighbor 20.1.1.3 local-as 200 no-prepend replace-as

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 remote-as 300
(config-router-af)#neighbor 3ffe:15:15:15:15::0 local-as 200
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 local-as 300
```

neighbor maximum-prefix

Use this command to set the number of prefixes that can be received from a neighbor.

Use the `no` parameter with this command to disable this function.

Note: For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) maximum-prefix <1-4294967295> <1-100|>
neighbor (A.B.C.D|X:X::X:X|WORD) maximum-prefix <1-4294967295> stop-update
neighbor (A.B.C.D|X:X::X:X|WORD) maximum-prefix <1-4294967295> warning-only
neighbor (A.B.C.D|X:X::X:X|WORD) maximum-prefix <1-4294967295> <1-100> warning-only
no neighbor (A.B.C.D|X:X::X:X) maximum-prefix <1-4294967295>
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-4294967295>`	Maximum number of prefixes accepted from this peer
`stop-update`	Stop installing routes when the maximum number of prefixes is exceeded.
`<1-100>`	Threshold value percent <1-100>
`warning-only`	Only give a warning message when the limit is exceeded. When this parameter is not specified and extra prefixes are received, the router ends the peering. A terminated peer remains down until the `clear ip bgp A.B.C.D` command is given.

Default

By default, neighbor maximum prefix is disabled

Command Mode

Address Family mode ad v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.72 maximum-prefix 1244 warning-only

(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 maximum-prefix 3
```

neighbor next-hop-self

Use this command to make the router the next hop for a BGP-speaking neighbor or peer group. This command allows a BGP router to change the nexthop information that is sent to the iBGP peer. The nexthop information is set to the IP address of the interface used to communicate with the neighbor.

Use the `no` parameter with this command to disable this feature.

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) next-hop-self
no neighbor (A.B.C.D|X:X::X:X|WORD) next-hop-self
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, next hop self is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.72 remote-as 100
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.72 next-hop-self

(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 next-hop-self
```

neighbor optional-as

Use this command to specify an AS (autonomous system) number to use with BGP dynamic peer-group.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor WORD optional-as <1-4294967295>
no neighbor WORD optional-as <1-4294967295>
```

Parameters

`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-4294967295>`	The range from which the optional AS number must be configured.

Default

By default, neighbor optional as is disabled

Command Mode

Router mode and Address Family VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor group1 optional-as 400
```

neighbor override-capability

Use this command to ignore received capabilities and use locally configured values.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) override-capability
no neighbor (A.B.C.D|X:X::X:X|WORD) override-capability
```

For BGP unnumbered mode:

```
neighbor WORD override-capability
no neighbor WORD override-capability
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, override-capability is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 override-capability
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 override-capability
```

neighbor passive

Use this command to make a BGP neighbor passive.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) passive
no neighbor (A.B.C.D|X:X::X:X|WORD) passive
```

For BGP unnumbered mode:

```
neighbor WORD passive
no neighbor WORD passive
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, neighbor passive is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 passive
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 passive
```

neighbor peer-group

Use this command to add a neighbor to an existing peer group. Neighbors with the same update policies are grouped into peer groups. This facilitates the updates of various policies, such as distribute and filter lists. The peer group is then configured easily with any of the neighbor commands. Any changes made to the peer group affect all members.

Note: See Usage below for when a peer group and a neighbor have conflicting attribute configurations.

To create a peer group, use the neighbor WORD peer-group command, and then use this command to add neighbors to the group.

Use the no parameter with this command to remove a neighbor from a named peer group.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X) peer-group WORD
no neighbor (A.B.C.D|X:X::X:X) peer-group WORD
```

For BGP unnumbered mode:

```
neighbor WORD peer-group WORD
no neighbor WORD peer-group WORD
```

Parameters

A.B.C.D	Address of the BGP neighbor in IPv4 format
X:X::X:X	Address of the BGP neighbor in IPv6 format
WORD	Peer group name

Default

No default value is specified

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor group1 peer-group
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

Note: For unnumbered peer we use peergroup key word to bind the unnumbered peer in group.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 peergroup ipi
```

Usage

When a *peer group* and a *peer* (neighbor) have conflicting attribute configurations the following rules apply:

- Outbound attribute configurations of a peer group *replace* peer member configurations of the same attributes when a peer becomes member of that peer group. Outbound attribute modifications to group members are not allowed.
- A peer group member's inbound attribute configurations take precedence over the peer group configuration.

These commands control outbound attribute updates:

- neighbor activate
- neighbor advertisement-interval
- neighbor as-origination-interval
- neighbor attribute-unchanged
- neighbor capability orf prefix-list
- neighbor distribute-list with an access-list number and the `out` parameter
- neighbor dont-capability-negotiate
- neighbor filter-list with the `out` parameter
- neighbor next-hop-self
- neighbor prefix-list with an access-list name and the `out` parameter
- neighbor remove-private-AS
- neighbor route-map with the `out` parameter
- neighbor route-reflector-client
- neighbor route-server-client
- neighbor send-community
- neighbor unsuppress-map

These commands control inbound attribute updates:

- neighbor allowas-in
- neighbor collide-established
- neighbor description
- neighbor distribute-list with an access-list number and the `in` parameter
- neighbor ebgp-multihop
- neighbor enforce-multihop
- neighbor filter-list with the `in` parameter
- neighbor g-shut
- neighbor g-shut-timer
- neighbor local-as
- neighbor maximum-prefix
- neighbor override-capability
- neighbor passive
- neighbor authentication-key
- neighbor port

- neighbor prefix-list with an access-list name and the `in` parameter
- neighbor remote-as
- neighbor restart-time
- neighbor route-map with the `in` parameter
- neighbor shutdown
- neighbor strict-capability-match
- neighbor update-source
- neighbor weight

BGP Peer Group Activation and Binding Guidelines

A BGP peer group can contain peer group members, including BGP IPv4, BGP IPv6, or BGP unnumbered peers. A peer group can consist of peer group members of different types; however, a restriction arises when the user attempts to activate the peer group across multiple address-families. Specifically, the user can only activate a BGP peer in a supported address family. This limitation can lead to functional issues and discrepancies in the show-running configuration.

To address this issue, limitations are implemented regarding the activation of the BGP peer group and the binding of peer group members.

- An error will occur if no peer group member is bound to the peer group during the activation attempt.
- A peer group that includes an IPv4 peer group member cannot also contain an IPv6 peer group member.
- Similarly, a peer group that includes an IPv6 peer group member cannot also contain an IPv4 peer group member.
- Additionally, a peer group that includes either an IPv4 or an IPv6 peer group member cannot contain an unnumbered peer.
- The BGP peer group type must remain unnumbered and cannot have IPv4 or IPv6 peers.

Applicability: OcNOS version 6.6.0

Note: Users must manually correct configurations after an upgrade or downgrade, as no scripts are available to facilitate this process.

neighbor WORD peer-group range

Use this command to create a dynamic peer group.

Use the no parameter with this command to remove a peer group.

Command Syntax

```
neighbor WORD peer-group range (A.B.C.D/M|X:X::X:X/M)
no neighbor WORD peer-group range (A.B.C.D/M|X:X::X:X/M)
```

Parameters

A.B.C.D/M	IP Prefix
X:X::X:X/M	IPv6 Prefix
WORD	Peer group name

Default

No default value is specified

Command Mode

Router mode and Address Family VRF mode.

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor group1 peer-group range 10.10.10.0/24
```

neighbor port

Use this command to set the BGP port number of a neighbor.

Use the `no` parameter with this command to remove a port number from a BGP neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) port <0-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) port
no neighbor (A.B.C.D|X:X::X:X|WORD) port <0-65535>
```

For BGP unnumbered mode:

```
neighbor WORD port <0-65535>
no neighbor WORD port
no neighbor WORD port <0-65535>
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<0-65535>`	Port number

Default

By default, neighbor port is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 port 643
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 port 500
```

neighbor prefix-list

Use this command to specify a prefix list for filtering BGP advertisements.

Filtering by prefix list matches the prefixes of routes with those listed in the prefix list. If there is a match, the route is used. An empty prefix list permits all prefixes. If a given prefix does not match any entries of a prefix list, the route is denied access. When multiple entries of a prefix list match a prefix, the entry with the smallest sequence number is considered to be a real match.

The router begins the search at the top of the prefix list, with the sequence number 1. Once a match or deny occurs, the router does not need to go through the rest of the prefix list. For efficiency the most common matches or denies are listed at the top. The neighbor distribute-list command is an alternative to this command and only one of them can be used to filter the same neighbor in any direction.

Use the `no` parameter with this command to remove an entry.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) prefix-list WORD (in|out)
no neighbor (A.B.C.D|X:X::X:X|WORD) prefix-list WORD (in|out)
```

For v4-unnumbered mode:

```
neighbor WORD prefix-list WORD (in|out)
no neighbor WORD prefix-list WORD (in|out)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Name of an access list
`in`	Apply access list to incoming advertisements
`out`	Apply access list to outgoing advertisements

Default

By default, neighbor prefix list is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ip prefix-list list1 deny 30.0.0.0/24
(config)#router bgp 12
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.10.10 prefix-list list1 in
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 prefix-list list2 out
```

neighbor remote-as

Use this command to establish a BGP peering relationship with a customer edge router.

The specified neighbor only exchanges unicast address prefixes, unless the neighbor is also activated using the neighbor activate command, which allows the exchange of other routing information.

Use the `no` parameter with this command to delete this peering.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) remote-as (<1-4294967295>|internal|external)
no neighbor (A.B.C.D|X:X::X:X|WORD) remote-as (<1-4294967295>|internal|external)
```

For BGP unnumbered mode:

```
neighbor WORD remote-as (internal|external)
no neighbor WORD remote-as (internal|external)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group. For an unnumbered interface (RFC 5549), specify an interface name configured with neighbor discovery or an IPv4 address with /31 or /30.
`<1-4294967295>`	Neighbor's autonomous system number (ASN) when extended capabilities are configured. If the specified ASN matches the ASN number specified in the router BGP global configuration, the neighbor is identified as internal. If the ASN does no match, the neighbor is identified as external to the local AS.
`internal`	iBGP peer.
`external`	eBGP peer.

Default

By default, neighbor remote is disabled

Command Mode

Router mode and Address Family-vrf mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 1.3.6.

Example

```
#configure terminal
(config)#router bgp 11
(config-router)#neighbor 10.10.0.73 remote-as 345
(config-router)#neighbor 11.11.0.74 remote-as 23456
```

Note: The last command in the example above should be used when the local speaker is OBGP and the neighbor is NBGP with a 4-octet ASN.

```
(config)#router bgp 100
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 10.10.0.1 remote-as 65000

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 remote-as 65000

(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 remote-as internal
(config-router-unnum)#show running bgp

router bgp 100
    !
    bgp unnumbered-mode
    neighbor eth1 remote-as internal
    neighbor eth2 remote-as external
    exit-unnumbered-mode
    !
```

neighbor remove-private-AS

Use this command to remove the private autonomous system number (ASN) from outbound updates. Private ASNs are not advertised to the Internet. This command is used with external BGP peers only. The router removes the ASNs only if the update includes private ASNs. If the update includes both private and public ASNs, the system treats it as an error.

Use the `no` parameter with this command to revert to default.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) remove-private-AS
no neighbor (A.B.C.D|X:X::X:X|WORD) remove-private-AS
```

For v4-unnumbered mode:

```
neighbor WORD remove-private-AS
no neighbor WORD remove-private-AS
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, neighbor remove private AS is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.63 remove-private-AS
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 remove-private-AS
```

© 2025 IP Infusion Inc. Proprietary

neighbor restart-time

Use this command to set a different restart-time than the global restart-time configured using the bgp graceful-restart command.

Use the `no` parameter with this command to restore the router to its default state.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) restart-time <1-3600>
no neighbor (A.B.C.D|X:X::X:X|WORD) restart-time <1-3600>
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-3600>`	The maximum time that a graceful-restart neighbor waits to come back up after a restart. Make sure that this value does not exceed the stalepath-time specified in router mode.

Default

By default, restart time is 90 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 3.3.3.3 restart-time 45
```

neighbor route-map

Use this command to apply a route map to incoming or outgoing routes. This command filters updates and modifies attributes. A route map is applied to inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates.

Use the `no` parameter with this command to a route map.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) route-map WORD (in|out)
no neighbor (A.B.C.D|X:X::X:X|WORD) route-map WORD (in|out)
```

For v4-unnumbered mode:

```
neighbor WORD route-map WORD (in|out)
no neighbor WORD route-map WORD (in|out)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Name of the route map
`in`	Apply access list to incoming advertisements
`out`	Apply access list to outgoing advertisements

Default

No default value is specified

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the configuration of the route map named `rmap2` and then the use of this map name in the `neighbor route-map` command.

```
#configure terminal
(config)#route-map rmap2 permit 6
(config-route-map)#match origin incomplete
(config-route-map)#set metric 100
(config-route-map)#exit
(config)#router bgp 12
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.10.10 route-map rmap2 in
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 route-map rmap1 in
```

neighbor route-reflector-client

Use this command to make the router a BGP route reflector and set a specified neighbor as its client.

Using route reflectors reduces the number of iBGP peers within an AS. An AS can have more than one route reflector. A route reflector treats other route reflectors as other iBGP speakers.

Use the `no` parameter with this command to indicate that the neighbor is not a client.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) route-reflector-client
no neighbor (A.B.C.D|X:X::X:X|WORD) route-reflector-client
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

No default value is specified

Command Mode

This command is supported in all Address Family modes.

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following configuration, Router1 is the route reflector for clients 3.3.3.3 and 2.2.2.2; it also has a non-client peer 6.6.6.6.

```
#configure terminal
(config)#router bgp 200
(config-router)#neighbor 3.3.3.3 remote-as 200
(config-router)#neighbor 2.2.2.2 remote-as 200
(config-router)#neighbor 6.6.6.6 remote-as 200
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 3.3.3.3 route-reflector-client
(config-router-af)#neighbor 2.2.2.2 route-reflector-client
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 route-reflector-client
```

For unnumbered peer, configs in L2vpn EVPN mode are configured under l2vpn-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family l2vpn evpn
(config-router-af)#neighbor eth1 route-reflector-client
```

neighbor route-server-client

Use this command to make a neighbor a route server client.

Use the `no` parameter with this command to remove the configuration of a neighbor as route server client.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) route-server-client
no neighbor (A.B.C.D|X:X::X:X|WORD) route-server-client
```

For v4-unnumbered mode:

```
neighbor WORD route-server-client
no neighbor WORD route-server-client
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

No default value is specified

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.72 route-server-client

#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#no neighbor 10.10.0.72 route-server-client
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 route-server-client
```

neighbor send-community

Use this command to send that a community attribute to a BGP neighbor.

The community attribute groups destinations in a certain community and applies routing decisions according to those communities. On receiving community attributes, the router reannounces them to the neighbor.

By default, both `standard` and `extended` community attributes are sent to a neighbor. To explicitly send only the `standard` or `extended` community attribute, run the bgp dampening command with the `standard` parameter before running this command.

Use the `no` parameter with this command to not announce community attributes to the neighbor. Use the `extended` and `no` parameters to remove extended communities. Specifying no other parameter with `no` removes standard communities only.

See also neighbor send-community in Chapter 2, *BGP Virtual Private Network Commands*.

Command Syntax

```
no neighbor (A.B.C.D|X:X::X:X|WORD) send-community
no neighbor (A.B.C.D|X:X::X:X|WORD) send-community (both|extended|standard)
neighbor (A.B.C.D|X:X::X:X|WORD) send-community
neighbor (A.B.C.D|X:X::X:X|WORD) send-community (both|extended|standard)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`both`	Send Standard and Extended Community attributes
`extended`	Send Extended Community attributes
`standard`	Send Standard Community attributes

Default

By default, both communities (standard and extended) are sent to every BGP neighbor.

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router bgp 100
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 10.10.10.1 remote-as 200
(config-router-af)#no neighbor 10.10.0.1 send-community extended
```

neighbor send-community large

Use this command to send that a community attribute to a BGP neighbor.

Use the `no` parameter with this command to not announce community attributes to the neighbor.

Command Syntax

```
neighbor A.B.C.D send-community large
no neighbor A.B.C.D send-community large
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format

Default

By default send-community large is enabled for peer

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 6.1.0.

Examples

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 1.1.1.1 activate
(config-router-af)#neighbor 1.1.1.1 send-community large
```

neighbor shutdown

Use this command to terminate active sessions for a specified neighbor and clear all related routing information.

If a peer group is specified, a large number of peering sessions might be terminated. The `show ip bgp summary` command displays a summary of BGP neighbors and their connections.

Use `description` option for operators to transmit a short free-form message to describe why a BGP session was shut down.

Use the `no` parameter with this command to re-enable a neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) shutdown  (description LINE|)
no neighbor (A.B.C.D|X:X::X:X|WORD) shutdown
```

Foe BGP unnumbered mode:

```
neighbor WORD shutdown (description LINE|)
no neighbor WORD shutdown
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`description`	Administratively shutdown communication msg to neighbor.

Default

By default, neighbor shutdown is disabbled

Command Mode

Router mode and Address Family-vrf mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3. The new version of the command with “description” option is introduced in OcNOS version 6.0.0.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.72 shutdown

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 shutdown description software
upgrade; back in 2 hours
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
```

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 shutdown
```

neighbor strict-capability-match

Use this command to close the BGP connection if a capability value does not match the remote peer.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) strict-capability-match
no neighbor (A.B.C.D|X:X::X:X|WORD) strict-capability-match
```

For BGP unnumbered mode:

```
neighbor WORD strict-capability-match
no neighbor WORD strict-capability-match
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, strict capability match is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 strict-capability-match
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 strict-capability-match
```

neighbor timers

Use this command to set the timers for a specific BGP neighbor.

Use the `no` parameter with this command to clear the timers for a BGP neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) timers <0-65535> <0-65535>
neighbor (A.B.C.D|X:X::X:X|WORD) timers connect <1-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) timers
no neighbor (A.B.C.D|X:X::X:X|WORD) timers connect
```

For BGP unnumbered mode:

```
neighbor WORD timers <0-65535> <0-65535>
neighbor WORD timers connect <1-65535>
no neighbor WORD timers
no neighbor WORD timers connect
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<0-65535>`	Keepalive interval. Keepalive messages are sent by a router to inform another router that the BGP connection between the two is still active. The keepalive interval is the period of time between each keepalive message sent by the router.
`<0-65535>`	Holdtime interval which is the time the router waits to receive a keepalive message. If the router does not receive a message in this period, the router declares the neighbor dead. The holdtime value should be at least 3 times the keepalive time.
`connect`	BGP connect timer
`<1-65535>`	Connect timer

Defaults

By default, keepalive timer value is 30 seconds

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 timers 60 230
```

```
(config-router)#neighbor 10.10.10.10 timers connect 10

(config-router)#no neighbor 10.10.10.10 timers
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 timers 40 120
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 timers connect 50
```

neighbor unsuppress-map

Use this command to selectively leak more-specific routes to a particular neighbor.

When the aggregate-address command is used with the `summary-only` option, the more-specific routes of the aggregate are suppressed to all neighbors. Use the `neighbor unsuppress-map` command to selectively leak more-specific routes to a particular neighbor.

Use the `no` parameter with this command to restore the setting to the default level.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) unsuppress-map WORD
no neighbor (A.B.C.D|X:X::X:X|WORD) unsuppress-map WORD
```

For v4-unnumbered mode:

```
neighbor WORD unsuppress-map WORD
no neighbor WORD unsuppress-map WORD
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Name of the route map used to select routes to unsuppress

Default

By default, neighbor unsuppress map is disabled

Command Mode

Address Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.0.70 unsuppress-map mymap
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 unsuppress-map map2
```

neighbor update-source

Use this command to allow internal BGP sessions to use any operating interface for TCP connections.

A loopback interface is most commonly used with this command. The use of loopback interface eliminates a dependency and BGP does not have to rely on the availability of a particular interface for making TCP connections.

Use the `no` parameter with this command to restore the interface assignment to the closest interface.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) update-source WORD
no neighbor (A.B.C.D|X:X::X:X|WORD) update-source
```

For BGP unnumbered mode:

```
neighbor WORD update-source WORD
no neighbor WORD update-source
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`WORD`	Interface name

Default

By default, neighbor update source is disabled

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.72 update-source myif
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 update-source myif
```

neighbor version

Use this command to accept only a particular BGP version.

By default, the system uses BGP version 4 and on request dynamically negotiates down to version 2. Using this command disables the router's version-negotiation capability and forces the router to use only a specified version with the neighbor.

Use the `no` parameter with this command to use the default version level of a neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) version 4
no neighbor (A.B.C.D|X:X::X:X|WORD) version
```

For BGP unnumbered mode:

```
neighbor WORD version 4
no neighbor WORD version
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`4`	BGP version number

Default

By default, the system uses BGP version 4 and on request dynamically negotiates down to version 2

Command Mode

Router mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#neighbor 10.10.10.10 version 4

(config)#router bgp 12
(config-router)#no neighbor 10.10.10.10 version
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 version 4
```

neighbor weight

Use this command to specify a weight value, per address-family, for all routes learned from a neighbor.

The route with the highest weight gets preference when the same prefix is learned from more than one peer. Unlike the local-preference attribute, the weight attribute is relevant only to the local router. The weights assigned using the `set weight` command override the weights assigned using this command.

Use this command in router mode to specify a weight value for all address families. Use this command in address family mode to specify a weight value per IPv4/IPv6/6PE address family,

When the weight is set for a peer group, all members of the peer group get the same weight. This command can also be used to assign a different weight to an individual peer-group member. When an individually-configured weight of a peer-group member is removed, its weight is reset to its peer group's weight.

Use the `no` parameter with this command to remove a weight assignment.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) weight <0-65535>
no neighbor (A.B.C.D|X:X::X:X|WORD) weight
no neighbor (A.B.C.D|X:X::X:X|WORD) weight <0-65535>
```

For v4-unnumbered mode:

```
neighbor WORD weight <0-65535>
no neighbor WORD weight
no neighbor WORD weight <0-65535>
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<0-65535>`	Weight value

Default

By default, neighbor weight value is 0

Command Mode

Router Address-Family mode and v4-unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 10.10.10.10 weight 60
```

```
(config-router-af)#no neighbor 10.10.10.10 weight 60
```

For unnumbered peer, IPv4 unicast mode commands are configured under v4-unnumbered-mode.

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor eth1 weight 44
```

neighbor WORD peer-group

Use this command to create a peer group.

Use the `no` parameter with this command to remove a peer group.

Command Syntax

```
neighbor WORD peer-group
no neighbor WORD peer-group
```

Parameters

`WORD`	Name of BGP peer group

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to create a peer group named `group1`.

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor group1 peer-group
```

network

Use this command to specify the networks to be advertised by the BGP routing process. A unicast network address without a mask is accepted if it falls into the natural boundary of its class. A class-boundary mask is derived if the address matches its natural class-boundary.

Use the `backdoor` parameter to specify a backdoor route to a BGP border router that will provide better information about the network. For data to be advertised by BGP, its routing table must include a route to the specified network. This command specifies the networks to be advertised. The network command works if the network being advertised is known to the router.

The `backdoor` parameter enables a route to be the preferred route even if it has a greater distance. A network that is specified as a backdoor network is dynamically assigned an administrative distance of 200 ensuring that IGP learned routes are preferred. If a backdoor network is not sourced by the local router, the network is learned from the external routers. If the route is learned from eBGP for a backdoor network, the distance is set to 20 or 200.

Use the `no` form of this command to remove a network route entry.

Command Syntax

```
network A.B.C.D (backdoor|)
network A.B.C.D/M (backdoor|)
network A.B.C.D mask A.B.C.D (backdoor|)
network A.B.C.D mask A.B.C.D route-map WORD (backdoor|)
network A.B.C.D route-map WORD (backdoor|)
network A.B.C.D/M route-map WORD (backdoor|)
no network A.B.C.D
no network A.B.C.D/M
```

Parameters

`A.B.C.D`	IP prefix <network>, for example, 35.0.0.0
`A.B.C.D/M`	IP prefix <network>/<length>, for example., 35.0.0.0/8
`backdoor`	BGP backdoor route
`routemap`	Route map used to modify the attributes
`WORD`	Name of the route map
`mask`	Network mask, for example, 255.255.0.0
`A.B.C.D`	Network mask, e.g., 255.255.0.0

Default

No default value is specified

Command Mode

Router Address-family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example illustrates a Class-A address configured as a network route. The natural Class-A network prefix mask length of 8 is internally derived, that is, 2.0.0.0/8.

```
(config)#router bgp 1
(config-router)#address-family ipv4 unicast
(config-router-af)#network 2.0.0.0
```

network synchronization

Use this command to enable IGP synchronization for BGP static network routes.

Use this no parameter with this command to disable synchronization of BGP static routes.

Command Syntax

```
network synchronization
no network synchronization
```

Parameters

None

Default

By default, network synchronization is disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example enables IGP synchronization of BGP static network routes in the router configuration mode.

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv4 unicast
(config-router-af)#network synchronization
```

The following example enables IGP synchronization of BGP static network routes in the IPv6 unicast address-family mode.

```
#configure terminal
(config)#router bgp 11
(config)#address-family ipv6 unicast
(config-router-af)#network synchronization
```

redistribute

Use this command to inject routes from one routing process into another. Redistribution is used by routing protocols to advertise routes that are learned by some other means, such as by another routing protocol or by static routes. Since all internal routes are dumped into BGP, careful filtering is applied to make sure that only routes to be advertised reach the internet, not everything. This command allows redistribution by injecting prefixes from one routing protocol into another routing protocol.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
redistribute [connected|isis|ospf|rip|static]
redistribute [connected|isis|ospf|rip|static] route-map WORD
no redistribute [connected|isis|ospf|rip|static]
no redistribute [connected|isis|ospf|rip|static] route-map
no redistribute [connected|isis|ospf|rip|static] route-map WORD
```

Parameters

`connected`	Redistribute connected routes
`isis`	Redistribute connected ISO IS-IS routes
`ospf`	Redistribute OSPFv2 routes
`rip`	Redistribute RIP routes
`static`	Redistribute static routes
`route-map`	Route map reference
`WORD`	Route map entries

Default

By default, redistribute is disabled

Command Mode

Address Family mode and Address Family-vrf mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the configuration of the route-map name `rmap1` and then the use of this map name in the `redistribute route-map` command.

```
#configure terminal
(config)#route-map rmap1 permit 1
(config-route-map)#match origin incomplete
(config-route-map)#set metric 100
(config-route-map)#exit
(config)#router bgp 12
(config-router)#address-family ipv4 unicast
(config-router-af)#redistribute ospf route-map rmap1
```

```
(config)#router bgp 100
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#redistribute static

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#redistribute static
```

restart bgp graceful

Use this command to enable a BGP-speaker router for graceful restart. This command stops the whole BGP process and makes OcNOS retain the BGP routes and mark them as stale. Receiving BGP speakers, retain and mark as stale all BGP routes received from the restarting speaker for all address families received in the Graceful Restart Capability exchange.

Command Syntax

```
restart bgp graceful
```

Parameters

None

Default

By default, bgp graceful is disabled

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#restart bgp graceful
#restart bgp graceful
% Warning : BGP process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
```

router bgp

Use this command to start a BGP process.

Use the `no` parameter with this command to disable an existing routing process.

Command Syntax

```
router bgp <1-4294967295>
no router bgp <1-4294967295>
```

Parameters

`<1-4294967295>`

Associate the routing process with this autonomous system number

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 12
(config-router)#
```

set large-community

Use this command to set the large community values in routing policies.

Command Syntax

```
set large-community (.AA:NN:MM|) (additive|)
no set large-community (.AA:NN:MM|) (additive|)
```

Parameters

`.AA:NN:MM`	Specifies the large community value with a maximum length of 255 characters.
`additive`	(Optional) Appends the specified large-community values to the existing large-community values in the route.

Default

The send-community large feature is enabled by default for peers.

Command Mode

Route-map mode

Applicability

Introduced in OcNOS version 6.1.0. Introduced new parameters and updated the maximum character length in OcNOS version 6.6.0.

Examples

Setting Multiple Large Communities

```
(config)#route-map R1 permit 10
(config-route-map)#set large-community 2:2:2 9:9:9 3:3:3
(config-route-map)#exit
```

The below configuration adds two large community attributes, `132:1:132` and `132:3:132`, to the matching routes. The `additive` option ensures these communities are added without replacing any existing ones.

```
(config)#route-map R1 permit 20
(config-route-map)#set large-community 132:1:132 132:3:132 additive
(config-route-map)#exit
```

Setting a Single Large Community

```
(config)#route-map R1 permit 15
(config-route-map)#set large-community 1:2:3
(config-route-map)#exit
```

snmp restart bgp

Use this command to restart SNMP in Border Gateway Protocol (BGP)

Command Syntax

```
snmp restart bgp
```

Parameters

None

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#snmp restart bgp
```

synchronization

Use this command to enable IGP synchronization of Internal BGP (iBGP) learned routes with the Internal Gateway Protocol (IGP) system in the router configuration mode or in the address-family configuration mode.

Synchronization is used when a BGP router should not advertise routes learned from iBGP neighbors, unless those routes are also present in an IGP (for example, OSPF). Synchronization may be enabled when all the routers in an autonomous system do not speak BGP, and the autonomous system is a transit for other autonomous systems.

The `no synchronization` command is used when BGP router can advertise routes learned from its iBGP neighbors without waiting for the IGP reachability to be present.

Command Syntax

```
synchronization
no synchronization
```

Parameters

None

Default

No default value is specified

Command Mode

Address Family modes

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example enables IGP synchronization of iBGP routes in Router mode.

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv4 unicast
(config-router-af)#synchronization
```

The following example enables IGP synchronization of iBGP routes in the IPv6-Unicast address family.

```
#configure terminal
(config)#router bgp 11
(config-router)#address-family ipv6 unicast
(config-router-af)#synchronization
```

timers bgp

Use this command to globally set or reset the keepalive and holdtime values for all the neighbors.

Use the `no` parameter with this command to reset timers to default value.

Command Syntax

```
timers bgp <0-65535> <0-65535>
no timers bgp
```

Parameters

`<0-65535>`	Frequency with which keepalive messages are sent to the neighbors
`<0-65535>`	Interval after which a neighbor is considered dead if keepalive messages are not received

Default

By default, keepalive timer value is 30 seconds

By default, holdtime value is 90 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#timers bgp 40 120
```

v4-unnumbered-mode

Use this command to configure the unnumbered peer with all the neighbor configs applicable in IPv4 unicast address family.

Use the `exit` parameter with this command to leave IPv4 BGP unnumbered mode.

Command Syntax

```
bgp v4-unnumbered-mode
exit-v4-unnumbered-mode
```

Parameters

None

Applicability

No default value is specified

Command Mode

Address-family IPv4 Unicast Mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#
```

bgp unnumbered-mode

Use this command to enter the unnumbered mode in router BGP mode allowing configuration of neighbor specific parameters. In this mode we allow creation of unnumbered peer and allow all the neighbor configuration applicable in router BGP mode.

Use no form of this command to exit unnumbered-mode and return to previous mode.

Note: ECMP with BGP unnumbered and numbered together is not supported.

Command Syntax

```
bgp unnumbered-mode
exit-unnumbered-mode
```

Parameters

None

Command Mode

Router BGP mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#
```

CHAPTER 2 BGP Virtual Private Network Commands

This chapter describes the BGP Virtual Private Network (VPN) configuration commands.

- address-family (see address-family in BGP Commands and address-family ipv6 in BGP4+ Commands sections.)
- bgp external-route-leak
- bgp inbound-route-filter
- exit-address-family (see exit-address-family in Chapter 1, *BGP Commands*)
- export map
- import map
- ip vrf
- neighbor activate (see neighbor activate in Chapter 1, *BGP Commands*)
- neighbor allow-ebgp-vpn
- neighbor allowas-in (see neighbor allowas-in in Chapter 1, *BGP Commands*)
- neighbor as-origination-interval (see neighbor as-origination-interval in Chapter 1, *BGP Commands*)
- neighbor as-override
- neighbor description (see neighbor description in Chapter 1, *BGP Commands*)
- neighbor remote-as (see neighbor remote-as in Chapter 1, *BGP Commands*)
- neighbor send-community (see neighbor send-community in Chapter 1, *BGP Commands*)
- neighbor shutdown (see neighbor shutdown in Chapter 1, *BGP Commands*)
- neighbor soo
- redistribute (see redistribute in Chapter 1, *BGP Commands*)
- rd (route distinguisher)
- route-target

© 2025 IP Infusion Inc. Proprietary

bgp external-route-leak

Use this command to control the external route leaking. An external imported route (those received from VPN neighbor and imported to an IP-VRF) is further leaked to another IP-VRF. The route-target exports the source VRF, matches the route-target, and imports the destination VRF.

Use the `no` parameter with this command to disable external route leaking.
External route leaking is only supported for BGP EVPN routes. It is not support for other types of VPN routes.

Command Syntax

```
bgp external-route-leak
```

```
no bgp external-route-leak
```

Parameters

None

Default

Enable

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.4

Examples

```
OcNOS#configure terminal
(config)#router bgp 100
(config-router)#no bgp external-route-leak
```

bgp inbound-route-filter

Use this command to control the filtering of received VPN routes with route-target extended community attributes. The inbound route filtering is applicable for both bgp inbound-route-filter and route-target import commands.

When a router is configured as EVPN Route-Reflector, it exchanges VRF routing information with a route distinguisher and route-target extended communities.

By default, OcNOS discards the received routes that does not match the local IP/MAC VRF's route-target import value. We can use command `no bgp inbound-route-filter` to override this behavior.

When the local box is acting as a EVPN route-reflector and not in the forwarding path, it may not be configured with an IP/MAC VRF terminations. In such case, `no bgp inbound-route-filter` is required to be configured to keep all the routes into RD (route-distinguisher) table.

Command Syntax

```
bgp inbound-route-filter
```

```
no bgp inbound-route-filter
```

Parameters

None

Default

By default, OcNOS does not import routing information that does not match the local IP or MAC VRF's route-target import value.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp inbound-route-filter
```

export map

This command assigns a route map to the VRF. This map is applied for routing information exported to another PE or VRF.

Use this command when an application requires finer control over the routes exported to another VRF or PE than provided by the import and export extended communities. You can filter routes that are eligible for export to another VRF or PE through the use of a route map. The route map can deny access to selected routes from a community that is on the export list.

Note: Only match rules (deny/permit) are applied for exported routes, set rules will not apply.

Use the `no` command to remove the map.

Command Syntax

```
export map WORD
no export map
```

Parameters

`WORD`	Route map

Default

No default value is specified

Command Mode

VRF mode

Applicability

This command was introduced in OcNOS version 4.1.

Examples

```
(config)#ip vrf myVRF
(config-vrf)#export map set-pref
(config-vrf)#
```

import map

This command assigns a route map to the VRF. This map is applied for routing information imported from another PE or VRF.

Use this command when an application requires finer control over the routes imported into a VRF than provided by the import and export extended communities. You can filter routes that are eligible for import into a VRF through the use of a route map.The route map can deny access to selected routes from a community that is on the import list.

Use the `no` option with this command to remove the map.

Command Syntax

```
import map WORD
no import map
```

Parameters

`WORD`	Route map

Default

No default value is specified

Command Mode

VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#ip vrf myVRF
(config-vrf)#import map set-pref
(config-vrf)#
```

ip vrf

Use this command to assign a VPN Routing Forwarding (VRF) instance.

Use the `no` option with this command to remove the VRF from the instance.

Command Syntax

```
ip vrf WORD
no ip vrf WORD
```

Parameter

`WORD`	Name of the VRF instance

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Command Example

```
(config)#ip vrf myVRF
(config-vrf)#
```

neighbor allow-ebgp-vpn

Use this command to allow an eBGP neighbor to be a VPN peer. By default, BGP VPN functionality is allowed only for iBGP peers.

Use the `no` parameter with this command to remove the configuration.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) allow-ebgp-vpn
no neighbor (A.B.C.D|X:X::X:X|WORD) allow-ebgp-vpn
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in IPv4 format
`X:X::X:X`	Address of the BGP neighbor in IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, BGP VPN functionality is allowed only for iBGP peers

Command Mode

Address Family-vpnv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router bgp 200
(config-router)#neighbor 66.66.66.66 remote-as 100
(config-router)#neighbor 66.66.66.66 update-source lo
(config-router-af)#neighbor 66.66.66.66 allow-ebgp-vpn
(config-router-af)#neighbor 66.66.66.66 activate
(config-router-af)#exit-address-family
```

neighbor as-override

Use this command to configure a provider edge (PE) router to override the autonomous system number (ASN) of a site with the ASN of a provider. BGP normally ignores routes from the same autonomous system. However, this command is used so that the Customer Edge (CE) routers router accepts and installs routes from the same autonomous system.

Typically, this command is used when CE routers have the same ASN in some or all sites. As per BGP requirement, a BGP speaker rejects a route that has the same ASN as itself in the `AS_PATH` attribute. Thus the CE routers having the same ASN do not accept routes from each other. Giving this command on the PE router removes the CE neighbor's ASN from the `AS_PATH` attribute allowing CE routers with the same ASN to accept routes from each other.

Use the `no` parameter with this command to remove VPN IPv4 prefixes from a specified router.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) as-override
no neighbor (A.B.C.D|X:X::X:X|WORD) as-override
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in IPv4 format
`X:X::X:X`	Address of the BGP neighbor in IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.

Default

By default, neighbor as override is disabled

Command Mode

Address Family-vrf mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 10.10.0.1 as-override

#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 as-override
```

neighbor send-community

Use this command to send the extended-community attribute to a customer edge router. In VPN, the route-distinguisher and route-target are encoded in BGP extended-community.

See also neighbor send-community in Chapter 1, *BGP Commands*.

Command Syntax

```
no neighbor (A.B.C.D|X:X::X:X|WORD) send-community
no neighbor (A.B.C.D|X:X::X:X|WORD) send-community (both|extended|standard)
neighbor (A.B.C.D|X:X::X:X|WORD) send-community
neighbor (A.B.C.D|X:X::X:X|WORD) send-community (both|extended|standard)
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`both`	Send standard and extended community attributes
`extended`	Send extended community attributes
`standard`	Send standard community attributes

Default

By default, both communities (standard and extended) are sent to every BGP neighbor.

Command Mode

Address Family Unicast mode and Address Family VRF mode.

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router bgp 100
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 10.10.10.1 remote-as 200
(config-router-af)#no neighbor 10.10.0.1 send-community extended
```

neighbor soo

Use this command to enable the site-of-origin (SOO) feature. If the customer AS is multi-homed to the ISP, this command ensures that the PE does not advertise the routes back to the same AS.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) soo AS:nn_or_IP:nn
no neighbor (A.B.C.D|X:X::X:X|WORD) soo
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in IPv4 format
`X:X::X:X`	Address of the BGP neighbor in IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`ASN:nn_or_IP-address:nn`	An AS number and an arbitrary number (for example, 100:1), or a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

By default, the site-of-origin (SOO) feature is disabled.

Command Mode

Address Family VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router bgp 100
(config-router)#address-family ipv4 vrf VRF_A
(config-router-af)#neighbor 1.1.1.1 remote-as 200
(config-router-af)#neighbor 10.10.0.1 soo 100:1
```

rd (route distinguisher)

Use this command to assign a route distinguisher (RD) for the VRF. The route distinguisher value must be a unique value on the router.

This command creates routing and forwarding tables and specifies the default RD for a VPN. The RD is added to the customer's IPv4 prefixes, changing them into globally unique VPN-IPv4 prefixes.

Use `no` form command to remove the RD configuration.

Note: RD configuration cannot be changed, it needs to be removed and added back with new value. When RD configuration is removed the RT configuration is also lost and needs to be reconfigured.

Command Syntax

```
rd ASN:nn_or_IP-address:nn
```

```
no rd ASN:nn_or_IP-address:nn
```

Parameters

`ASN:nn_or_IP-address:nn`

AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

No default value is specified

Command Mode

VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#ip vrf VRF_A
(config-vrf)#rd 100:1
```

route-target

Use this command to add a list of import and export route-target extended communities to the VRF.

This command creates lists of import and export route-target extended communities for the VRF. It specifies a target VPN extended community. Execute the command once for each community. All routes with the specific route-target extended community are imported into all VRFs with the same extended community as an import route-target.

Use the `no` parameter with this command to delete a route target.

Command Syntax

```
route-target (import|export|both) ASN:nn_or_IP-address:nn
no route-target (import|export|both) ASN:nn_or_IP-address:nn
```

Parameters

`import`	Import routing information
`export`	Export routing information
`both`	Import and export routing information
`ASN:nn_or_IP-address:nn`	AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

No default value is specified

Command Mode

VRF mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#ip vrf VRF_A
(config-vrf)#route-target both 100:10

(config)#ip vrf VRF_A
(config-vrf)#route-target import 100:20
```

CHAPTER 3 BGP Show Commands

This chapter describes the BGP show commands.

- show bgp
- show bgp A.B.C.D
- show bgp A.B.C.D/M
- show bgp client
- show bgp community
- show bgp community-list
- show bgp dampening dampened-paths
- show bgp dampening flap-statistics
- show bgp dampening parameters
- show bgp filter-list
- show bgp inconsistent-as
- show bgp ipv6
- show bgp neighbors
- show bgp neighbors
- show bgp neighbors advertised-routes
- show bgp neighbors received prefix-filter
- show bgp neighbors received-routes
- show bgp neighbors routes
- show bgp nexthop-tracking
- show bgp nexthop-tree-details
- show bgp paths
- show bgp prefix-list
- show bgp quote-regexp
- show bgp regexp
- show bgp route-map
- show bgp statistics
- show bgp summary
- show bgp X:X::X:X
- show bgp X:X::X:X/M longer prefixes
- show debugging bgp
- show ip bgp
- show ip bgp cidr-only
- show ip bgp community-info
- show ip bgp peer-group
- show ip bgp peer-group vrf all

© 2025 IP Infusion Inc. Proprietary

- show ip bgp rtfilter all
- show ip bgp scan
- show ip extcommunity-list
- show ip extcommunity-list
- show ip protocols
- show ip vrf
- show running-config as-path access-list
- show running-config bgp
- show running-config community-list
- show running-config vrf
- show running-config ip vrf
- show running-config mac vrf
- show running-config vrf word

show bgp

Use this command to display the status of BGP routes.

Command Syntax

```
show bgp
show bgp (ipv6)
show bgp (ipv4|ipv6) (unicast|multicast)
show ip bgp
show ip bgp ipv4 (unicast|multicast)
show bgp (vrf (VRFNAME|all|default))
show ip bgp (vrf (VRFNAME|all|default))
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip bgp
BGP table version is 1, local router ID is 12.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop      Metric    LocPrf   Weight   Path
*> 1.1.1.1/32   0.0.0.0          0         100      32768      i

Total number of prefixes 1
```

show bgp A.B.C.D

Use this command to display BGP route information for a network.

Command Syntax

```
show bgp (ipv4) (unicast|multicast) A.B.C.D
show ip bgp A.B.C.D
show ip bgp ipv4 (unicast|multicast) A.B.C.D
```

Parameters

`ipv4`	IPv4 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`A.B.C.D`	IP prefix (network), for example, 35.0.0.0

Command Mode

Privileged Exec and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
OcNOS#show ip bgp 80.80.80.0
BGP routing table entry for 80.80.80.0/24
Paths: (1 available, no best path)
 Not advertised to any peer
 300
   15.15.15.1 (inaccessible) from 11.11.11.2 (15.15.15.2)
     Origin incomplete, metric 0, localpref 100, valid, internal
     rx path_id: 0      tx path_id: -1
     Last update: Wed May 11 15:22:36 2016
```

Table 1 explains the output fields.

Table 1: show ip bgp output details

Field	Description
Paths	The paths listed in the routing table, along with path information, and whether the path are being advertised.
Metric	If shown, the value of the inter-autonomous system metric.
LocalPref	Local preference value as set with the set local-preference route-map configuration command.
rx path_id	Autonomous system receive path to the source network. There can be one entry in this field for each autonomous system in the path.

Table 1: show ip bgp output details

Field	Description
tx path_id	Autonomous system transmit path to the destination network. There can be one entry in this field for each autonomous system in the path.
Last update	Last time since the neighbor transitioned to or from the established state.

show bgp A.B.C.D/M

Use this command to display BGP route information for a network prefix.

Command Syntax

```
show bgp ipv4 (unicast|multicast) A.B.C.D/M
show ip bgp A.B.C.D/M
show ip bgp ipv4 (unicast|multicast) A.B.C.D/M
show ip bgp A.B.C.D/M (vrf (VRFNAME|all|default))
show ip bgp A.B.C.D/M longer-prefixes
show ip bgp ipv4 (unicast|multicast) A.B.C.D/M longer-prefixes
show ip bgp A.B.C.D/M longer-prefixes (vrf (VRFNAME|all|default))
```

Parameters

`ipv4`	IPv4 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`A.B.C.D/M`	IP prefix (network/length), for example, 35.0.0.0/8
`longer-prefixes`	Display route and more specific routes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp ipv4 unicast 22.22.22.22/32
BGP routing table entry for 22.22.22.22/32
Paths: (2 available, best #1, table Default-IP-Routing)
  Not advertised to any peer
  Local
    11.11.11.2 from 11.11.11.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Last update: Fri Feb 16 09:46:54 2001

  Local
    11.11.12.2 from 11.11.12.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal
      Last update: Fri Feb 16 09:46:52 2001
```

show bgp client

Use this command to display BGP client information.

Command Syntax

```
show bgp client
```

Parameters

None

Command Mode

Privileged Exec and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp client
R1#sh bgp client
BGP client ID: 11
 PIM, socket 10
  Service: AS number service, Route Service
  Messsage received 1, sent 1
  Connection time: Tue May 14 03:11:01 2019
  Last message read: Service Request
  Last message write: Service Reply
R1#
```

show bgp community

Use this command to display BGP routes that match a community.

Command Syntax

```
show bgp ip (unicast|multicast) community
show bgp ip (unicast|multicast) community (vrf (VRFNAME|all|default))
show bgp ip (unicast|multicast) community [AA:NN|local-AS|no-advertise|no-export]
  (exact-match|)
show ip bgp community
show ip bgp community [AA:NN|local-AS|no-advertise|no-export|internet] (exact-
  match|)
show ip bgp community (vrf (VRFNAME|all|default))
show ip bgp community [AA:NN|local-AS|no-advertise|no-export|internet] (exact-
  match|)
show ip bgp community-list WORD (exact-match|) (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) community
show ip bgp ipv4 (unicast|multicast) community (vrf (VRFNAME|all|default))
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF
`AA:NN`	Community number
`local-AS`	Do not send outside local AS (well-known community)
`no-advertise`	Do not advertise to any peer (well-known community)
`no-export`	Do not export to next AS (well-known community)
`internet`	Internet community (well-known community)
`exact-match`	Exact match of the communities

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp community
BGP table version is 1, local router ID is 102.67.98.95
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network                Next H     Metric    LocPrf     Weight    Path
*> 2001:43f8:bb1::/64   ::          0        100       32768      ?

Total number of prefixes 1
```

show bgp community-list

Use this command to display BGP routes that match a community list.

Command Syntax

```
show bgp community-list WORD (exact-match|)
show bgp ip (unicast|multicast) community-list WORD (exact-match|)
show bgp ip (unicast|multicast) community-list WORD (exact-match|)
show bgp (ipv4|ipv6) (unicast|multicast) community-list WORD (exact-match|)
show bgp (ipv6) community-list WORD (exact-match|)
show ip bgp community-list WORD (exact-match|)
show ip bgp ipv4 (unicast|multicast) community-list WORD (exact-match|)
```

Parameters

`WORD`	Community list name
`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`exact-match`	Only routes that exactly match the community

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp community-list Originate
Named Community expanded list Originate
    permit 37721:1[0-2]0[05]$
    permit 37721:1[0-2]0[05]
```

show bgp dampening dampened-paths

Use this command to display detailed information about paths suppressed due to dampening.

Command Syntax

```
show bgp dampening dampened-paths
show bgp (ipv4|ipv6) (unicast|multicast) dampening dampened-paths
show bgp (ipv6) dampening dampened-paths
show ip bgp dampening dampened-paths
show ip bgp dampening dampened-paths (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) dampening dampened-paths
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp dampening dampened-paths
```

show bgp dampening flap-statistics

Use this command to display BGP dampening flap statistics.

Command Syntax

```
show bgp dampening flap-statistics
show bgp (ipv4|ipv6) (unicast|multicast) dampening flap-statistics
show bgp (ipv6) dampening flap-statistics
show ip bgp dampening flap-statistics
show ip bgp dampening flap-statistics (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) dampening flap-statistics
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This sample output shows that the internal route (`i`), has flapped 3 times and is now categorized as history (`h`).

```
#show ip bgp dampening flap-statistics
BGP table version is 1, local router ID is 30.30.30.77
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
hi1.1.1.0/24     10.100.0.62       3  00:01:20       i
```

Header

```
BGP table version is 1, local router ID is 30.30.30.77
```

- BGP table version
- BGP router ID is `30.30.30.77`

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S stale
```

Table 1 shows the status codes displayed at the start of a route entry.

Table 2: status details

Status Code	Description	Comments
s	suppressed	Whether the route is suppressed and is not advertised to neighbors.
d	damped	When the penalty of a flapping route exceeds the suppress limit, the route is damped and remains in a withdrawn state until its penalty decreases below the reuse limit.
h	history	When the penalty of a flapping route does not exceed the suppress limit, the route is not damped and BGP maintains a history of the flapping route.
p	stale	When the BGP neighbor from which a route is learned is in graceful restart, the route is retained in the BGP routing table, but marked as stale .
*	valid	Whether the route is valid. When a route is not suppressed, damped, or present in the history, it is valid.
>	best	The selected route to be installed in the kernel routing table.
i	internal	The prefix was learned from an iBGP peer.

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Table 3 shows the codes at the end of each route entry that indicate where the route originated.

Table 3: origin details

Origin Code	Description	Comments
i	IGP	The route is from an Interior Gateway Protocol.
e	EGP	The route is from an Exterior Gateway Protocol.
?	incomplete	Origin not known. Typically, these are routes redistributed from an IGP.

Table 4 explains the output fields.

Table 4: show bgp dampening flap-statistics output details

Field	Description
Network	Internet address of a network.
From	IP address of the advertising peer.
Flaps	Number of times this route has failed and returned (flapped).

Table 4: show bgp dampening flap-statistics output details

Field	Description
Duration	Elapsed time since the first penalty points were assessed.
Reuse	The amount of time remaining until this route will be un-suppressed and can be used again.

show bgp dampening parameters

Use this command to display the BGP dampening parameters.

Command Syntax

```
show bgp dampening parameters
show bgp (ipv4|ipv6) (unicast|multicast) dampening parameters
show bgp (ipv6) dampening parameters
show ip bgp dampening parameters
show ip bgp ipv4 (unicast|multicast) dampening parameters (vrf
  (VRFNAME|all|default))
show ip bgp dampening parameters (vrf (VRFNAME|all|default))
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip bgp dampening parameters
 dampening 5 750 2000 60 15
 Dampening Control Block(s):
  Reachability Half-Life time    : 5 min
  Reuse penalty                  : 750
  Suppress penalty               : 2000
  Max suppress time              : 60 min
  Un-reachability Half-Life time : 15 min
  Max penalty (ceil)             : 11999
  Min penalty (floor)            : 375
#
```

Table 5 explains the output fields.

Table 5: show ip bgp dampening parameters output details

Field	Description
Dampening Control Block(s)	Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route.
Reachability Half-Life time	Number of minutes after which an arbitrary value is halved if a route stays stable.
Reuse penalty	Reuse threshold—Arbitrary value below which a suppressed route can be used again.
Suppress penalty	Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements.
Max suppress time	Maximum hold-down time for a route, in minutes.
Un-reachability Half-Life time	Number of minutes after which an arbitrary value is not halved if a route stays stable.
Max penalty (ceil)	Maximum penalty corresponds to the time it would take max-suppress to decay and reach the reuse level (ceil).
Min penalty (floor)	Maximum penalty corresponds to the time it would take max-suppress to decay and reach the reuse level (floor).

show bgp filter-list

Use this command to display routes that match a regular expression filter list.

Command Syntax

```
show bgp filter-list WORD
show bgp (ipv4|ipv6) (unicast|multicast) filter-list WORD
show bgp (ipv6) filter-list WORD
show ip bgp filter-list WORD
show ip bgp ipv4 (unicast|multicast) filter-list WORD
show ip bgp filter-list WORD (exact-match)
show ip bgp filter-list WORD (exact-match) (vrf (VRFNAME|all|default))
show ip bgp filter-list WORD (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) filter-list WORD (exact-match)(vrf
  (VRFNAME|all|default))
```

Parameters

`WORD`	Regular-expression filter list
`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`exact-match`	Exact match of the filter list
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF
`vrf`	VPN Routing/Forwarding instance name

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp filter-list bgp-local-only
BGP table version is 1, local router ID is 102.67.98.95
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network               Next Hop        Metric    LocPrf      Weight   Path
*> 2001:43f8:bb1::/6     ::              0        100         32768     ?
```

```
Total number of prefixes 1
```

show bgp inconsistent-as

Use this command to display routes with inconsistent AS paths.

Command Syntax

```
show bgp inconsistent-as
show bgp (ipv4|ipv6) (unicast|multicast) inconsistent-as
show bgp (ipv6) inconsistent-as
show ip bgp inconsistent-as
show ip bgp ipv4 inconsistent-as
show ip bgp ipv4 (unicast|multicast) inconsistent-as
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp inconsistent-as

BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop              Metric         LocPrf   Weight   Path
*> 10.1.0.0   172.29.232.55            0             0         100      ?
```

show bgp ipv6

Use this command to display the BGP routing table.

Command Syntax

```
show bgp ipv6 (unicast|multicast|labeled|)
show bgp ipv6 (unicast|multicast|labeled|) X:X::X:X/M
```

Parameters

`multicast`	IPv6 multicast address prefixes
`unicast`	IPv6 unicast address prefixes
`labeled`	Labeled IPv6 routes
`X:X::X:X/M`	IPv6 prefix network/length, such as 3ffe:a::/64

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example: iBGP and eBGP Routes

This example shows routes learned from both iBGP and eBGP.

```
#show bgp ipv6
BGP table version is 0, local router ID is 10.100.0.77
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal S stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network                                  Metric LocPrf Weight  Path
*> 2001:58::/32                                                0  20 ?
     fe80::202:b3ff:fec8:9fdb
*> 2002:58::/32                                                0  20 i
     fe80::202:b3ff:fec8:9fdb
*>i2003:58::/32                                  100           0  i
     fe80::208:a1ff:fe16:797d
```

Header

```
BGP table version is 0, local router ID is 10.100.0.77
```

- BGP table version
- BGP router ID is `10.100.0.77`

```
Status codes: s suppressed, d damped, h history, p stale, * valid, > best, i - internal
```

Table 6 shows the status codes displayed at the start of a route entry.

Table 6: status details

Status Code	Description	Comments
s	suppressed	Whether the route is suppressed and is not advertised to neighbors.
d	damped	When the penalty of a flapping route exceeds the suppress limit, the route is damped and remains in a withdrawn state until its penalty decreases below the reuse limit.
h	history	When the penalty of a flapping route does not exceed the suppress limit, the route is not damped and BGP maintains a history of the flapping route.
p	stale	When the BGP neighbor from which a route is learned is in graceful restart, the route is retained in the BGP routing table, but marked as stale .
*	valid	Whether the route is valid. When a route is not suppressed, damped, or present in the history, it is valid.
>	best	The selected route to be installed in the kernel routing table.
i	internal	The prefix was learned from an iBGP peer.

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Table 7 shows the codes at the end of each route entry that indicate where the route originated.

Table 7: origin codes

Origin Code	Description	Comments
i	IGP	The route is from an Interior Gateway Protocol.
e	EGP	The route is from an Exterior Gateway Protocol.
?	incomplete	Origin not known. Typically, these are routes redistributed from an IGP.

Route Entry Examples

```
*> 2002:58::/32     fe80::202:b3ff:fec8:9fdb                              0 20 i
```

- This route entry shows that this route is learned from eBGP.
- The origin code “i” means that the prefix is added by the network statement at originating AS.
- The path 20 indicates that the prefix advertisement originated from AS20.
- The administrative weight parameter applies only to routes within an individual router.
- Since this route was learned from a peer, it has a default weight of 0. All routes generated by the local router have a weight of 32,768.

```
*> 2001:58::/32      fe80::202:b3ff:fec8:9fdb                              0 20 ?
```

- This route entry shows that the prefix is learnt from eBGP.The origin code i indicates that the prefix is added by network statement at originating AS.The path attribute 20 indicates that the route advertisement originated from AS20.The administrative weight parameter applies only to routes within an individual router. Since this

route was learned from a peer, it has a default weight of 0. All routes generated by the local router have a weight of 32,768.The origin code “?” means the route was learned through redistribution.

```
*>i2003:58::/32     fe80::208:a1ff:fe16:797d                    100     0 i
```

- The status code “i” means that the route was learned through iBGP. The Local Preference attribute of the route, which is used only with the local AS, is set to 100 (the default value).

Example: IPv6 Prefix Routes

This example shows labeled routes for a given IPv6 prefix:

```
#show bgp ipv6 labeled 3ffe:a::/64
           BGP routing table entry for 3ffe:a::/64
           Paths: (1 available, best #1, table Default-IP-Routing-Table)
           Not advertised to any peer
           Local
           ::ffff:114:1414 from 20.20.20.1 (92.92.92.92)
          Origin incomplete metric 0, localpref 100, label    5420,
           valid, internal, best
           Last update: Mon May 26 17:48:18 2008
```

Table 8 explains the output fields.

Table 8: show bgp ipv6 output details

Field	Description
Paths	The paths listed in the routing table, along with path information, and whether the path are being advertised.
Metric	If shown, the value of the inter-autonomous system metric.
Local Pref	Local preference value as set with the set local-preference route-map configuration command.
rx path_id	Autonomous system receive path to the source network. There can be one entry in this field for each autonomous system in the path.
tx path_id	Autonomous system transmit path to the destination network. There can be one entry in this field for each autonomous system in the path.
Last update	Last time since the neighbor transitioned to or from the established state.

show bgp ipv6 peer-group

Use this command to list the BGP peer group information in ipv6 unicast.

Command Syntax

```
show bgp ipv6 peer-group (WORD|)
```

Parameters

WORD	Name of the peer group

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#show bgp ipv6 peer-group p2
BGP dynamic peer-group is p2, IBGP, remote AS 42949672
  BGP dynamic peer-group p2 listen range group members:
  5ffe:3::/64
  BGP version 4
Minimum time between advertisement runs is 5 seconds
For address family: IPv6 Unicast
  Peer-group member:
  *5ffe:3::1
  Index 2, Offset 0, Mask 0x4
  0 accepted prefixes, 0 announced prefixes
#
```

show bgp neighbors

Use this command to display information about BGP neighbor connections.

Command Syntax

```
show bgp neighbors
show bgp neighbors (A.B.C.D|X:X::X:X|WORD)
show bgp ipv6 neighbors
show ip bgp ipv4 (unicast|multicast) neighbors
show ip bgp neighbors
show ip bgp neighbors (A.B.C.D|X:X::X:X)(advertised-routes|)
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X|WORD)
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) (hold-time|keepalive-
  interval|connection-retrytime)
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) (sent-msgs|rcvd-
  msgs|notification|update|open|keepalive)
```

Parameters

`ipv4`	IPv4 neighbors
`ipv6`	IPv6 neighbors
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`A.B.C.D`	IPv4 neighbor
`X:X::X:X`	IPv6 neighbor
`WORD`	Interface name
`advertised-routes`	Routes advertised to a BGP neighbor
`hold-time`	Hold time
`keepalive-interval`	Keepalive interval
`connection-retrytime`	Connection retry time
`sent-msgs`	Sent packets
`rcvd-msgs`	Received packets
`notification`	Notification messages
`update`	Update messages
`open`	Open messages
`keepalive`	Keepalive messages

Command Mode

Privileged Exec and Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp neighbors
BGP neighbor is 2.2.2.2, remote AS 200, local AS 200, internal link
 Member of peer-group myPeer for session parameters
  BGP version 4, remote router ID 10.12.7.155
  BGP state = Established, up for 00:04:55
  Last read 00:04:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 IPv4
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  myPeer peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 2.2.2.1, Local port: 33865
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 2.2.2.1
Nexthop global: 1111::1
Nexthop local: fe80::a00:27ff:fecc:47a6
BGP connection: non shared network
Last Reset: 00:32:48, due to BGP Notification sent
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

BGP Neighbor Fields

Table 9 explains the output fields.

Table 9: show bgp neighbor output details

Field	Description
BGP neighbor	BGP session information for the neighbor with the ip-address argument.
remote AS	Remote Autonomous system used to exchange exterior routing information between neighboring ASs.

Table 9: show bgp neighbor output details

Field	Description
local AS	Local Autonomous system used to exchange internal routing information within AS.
internal link external link	Internal link is used to forward route advertisements received from an external BGP router through the internal network (in the same AS). External link is used for exchanging routing information between Autonomous Systems (AS) and routing traffic across the Internet (eBGP neighbor).
member of peer-group	Peer group information for the peer group specified with the peer-group argument.
BGP version	Negotiated BGP version for this session.
remote router ID	IP address of the neighbor. BGP uses the highest loopback address as the Router ID. If no loopback interface is configured, BGP uses the highest configured IP address on a system.
BGP state	Session state as explained in Table 14. The exchange of routing information begins between peers only after the neighbor session is in an Established state.
up for	Time that the underlying TCP connection has been up.
last read	Time since BGP last received a message from this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving messages. The maximum time that can elapse between successive messages from this neighbor is 180 seconds. If no message is received for 180 seconds, this neighbor will be declared dead.
last write	Time since BGP last sent a message to this neighbor.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor. The time interval between successive keepalive messages is `60` seconds. Typically, the hold time value is set to three times the keepalive interval.
neighbor capabilities	BGP capabilities advertised and received from this neighbor. “Advertised and received” is displayed when a capability is successfully exchanged between two routers.
received	Total number of received messages. notifications: Number of notification (error) messages received. in queue: Number of messages in the input queue
sent	Total number of sent messages. notifications: Number of notification (error) messages sent. in queue: Number of messages in the output queue
route refresh request	Number of route refresh request messages sent and received.
minimum time between advertisement runs.	The minimum time gap, in seconds, between successive route updates sent to the neighbor. Generally, a jitter (of 25%) is applied to this time interval, which means that if the time between advertisements is configured as 30, successive advertisements can have a time gap of as low as 22.5 (after applying a 25% jitter to the 30 seconds, which is 7.5 seconds).
for address family	The peers have exchanged address family capability.
BGP table version	For each of the address families agreed upon, BGP maintains a separate table.
neighbor version	Tracks prefixes that have been sent and those that need to be sent.

Table 9: show bgp neighbor output details

Field	Description
connections established	The number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other. “Dropped” means the number of time the connection has failed or gone down.
local host foreign host	Local host is the IP address and the port number of the local system used for the peering session. Foreign host is the IP address and the port of the neighbor. BGP always uses the TCP port number 179 for the peer originating the session.
nexthop	The IP address of the next hop used to reach the neighbor. eBGP or iBGP peers do not need to be directly connected. Peering sessions can be set up across multiple hops. If the neighbors are directly connected, the IP address of the local system is listed as the next hop.
nexthop global	The global IPv6 address of the next hop
nexthop local	The link-local IPv6 address of the next hop
non shared network	The peering session is running on a non shared network.
last reset	Time since this peering session was last reset. The reason for the reset is displayed on this line.
notification error message	Last error message sent.

show bgp neighbors advertised-routes

Use this command to display the routes advertised to a BGP neighbor.

Command Syntax

```
show bgp neighbors (A.B.C.D|X:X::X:X|WORD) advertised-routes
show ip bgp neighbors (A.B.C.D|X:X::X:X) advertised-routes (vrf
  (VRFNAME|all|default))
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) advertised-routes
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) routes advertised
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X|WORD) advertised-
  routes
```

Parameters

`A.B.C.D`	IPv4 neighbor
`X:X::X:X`	IPv6 neighbor
`WORD`	Interface name
`ipv4`	IPv4 addresses
`multicast`	Multicast prefixes
`unicast`	Unicast prefixes
`vrf`	VPN Routing/Forwarding instance name

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp neighbors 11.11.11.2 advertised-routes
BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network              Next Hop            Metric    LocPrf       Weight   Path
*>i 33.33.33.33/32   11.11.11.1           0         100         32768     i

Total number of prefixes 1
```

show bgp neighbors received prefix-filter

Use this command to display the prefix list filter.

Command Syntax

```
show bgp neighbors (A.B.C.D|X:X::X:X|WORD) advertised-routes
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) received prefix-filter
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X) received prefix-
  filter
```

Parameters

`A.B.C.D`	IPv4 address
`X:X::X:X`	IPv6 address
`WORD`	Interface name
`ipv4`	IPv4 addresses
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp neighbors 100.1.1.1 received prefix-filter
Address family: IPv4 Unicast
ip prefix-list 100.1.1.1.1.1: 1 entries
   seq 5 permit 1.1.1.1/32
```

show bgp neighbors received-routes

Use this command to display the received routes from a neighbor.

To display all the received routes from a neighbor, perform a BGP soft reconfigure first.

Command Syntax

```
show bgp neighbors (A.B.C.D|X:X::X:X|WORD) received-routes
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X) received-routes
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) received-routes
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) routes received
```

Parameters

`A.B.C.D`	IPv4 address
`X:X::X:X`	IPv6 address
`WORD`	Interface name
`ipv4`	IPv4 addresses
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp neighbors 11.11.11.1 received-routes
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network              Next Hop       Metric    LocPrf       Weight   Path
*>i 33.33.33.33/32   11.11.11.1      0         100          0        i

Total number of prefixes 1
```

show bgp neighbors routes

Use this command to display all accepted routes learned from neighbors.

Command Syntax

```
show bgp neighbors (A.B.C.D|X:X::X:X|WORD) routes
show ip bgp neighbors (A.B.C.D|X:X::X:X|WORD) routes
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X|WORD) routes
```

Parameters

`A.B.C.D`	IPv4 address
`X:X::X:X`	IPv6 address
`WORD`	Interface name
`ipv4`	IPv4 addresses
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following output displays detailed information about the neighbor.

```
#show bgp neighbors 10.10.10.2 routes
BGP neighbor is fe80::203:47ff:feb0:d72b, remote AS 10, local AS 10, internal
link
  BGP version 4, remote router ID 10.10.10.50
  BGP state = Established, up for 00:02:01
  Last read 00:00:01, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 3 messages, 0 notifications, 0 in queue
  Sent 5 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
```

Table 10 explains the output fields.

Table 10: show bgp neighbors routes output details

Field	Description
BGP neighbor	Neighbor identifier along with the local and remote Autonomous System numbers.
BGP version	The version of BGP being used by the neighbor device, along with the remote router ID number.
BGP state	The current state of the neighbor connection and length of time within the current state. Possible states are: Idle, Connect, Active, and Established.
Last read	The amount of time in Hours**:** Minutes**:** Seconds since this device last checked the Hold Time parameters.
hold time	The amount of time this device waits for a Keepalive or Update message before the BGP connection is closed.
Keepalive interval	KEEPALIVE messages are sent periodically to ensure that the connection is live.
Neighbor capabilities	New or optional parameters called “Capabilities.” Provides a graceful way to advertise new or unique options without causing peering to terminate. The capabilities are communicated in TLV fields. (see RFC 3392). In the example output above, the following capabilities were advertised by the neighbor and were received and understood by this device: • Route refresh • Address family IPv4 Unicast • Address family IPv6 Unicast
Route refresh	This helps to identify that and synchronize the peers without a hard reset.
For address family	Address Family Identifier (AFI) IPv4 Unicast.
Received messages	Information received from the neighbors.
notifications	Passes information to a router about sessions.
in que	Number of messages currently residing in the queue awaiting action.
Route refresh request	Information received and sent.
Minimum time between advertisement runs	Interval between exchange of messages.
For address family: IPv4 Unicast	The following three values are subordinate to the IPv4 Unicast Address Family.
Community attributes sent to this neighbor	Both the standard and the extended community information has been sent to the neighbor.
accepted prefixes	Configure a limit to the number of prefixes that can be accepted in a BGP peer session.
announced prefixes	A prefix announced in BGP consists of the IPV4 or IPV6 address block being announced.

show bgp nexthop-tracking

Use this command to display BGP nexthop-tracking status.

Command Syntax

```
show bgp nexthop-tracking
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp nexthop-tracking

Configured NHT: ENABLED
NHT Delay time-interval : 6
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 100, router-id 4.4.4.40
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
Total number of IPV4 nexthops : 0
Total number of IPV6 nexthops : 0

BGP VRF: VRF_A VRF_ID 2
BGP Instance: (Default), AS: 100, router-id 4.4.4.40
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
Total number of IPV4 nexthops : 0
Total number of IPV6 nexthops : 0
```

Table 11 explains the output fields.

Table 11: show bgp nexthop-tracking output details

Field	Description
Configured NHT	Whether Next Hop Tracking (NHT) is enabled or disabled.
NHT Delay time-interval	A delay timer that indicates how long this device waits before checking its RIB for changes.

Table 11: show bgp nexthop-tracking output details

Field	Description
BGP VRF	Name and ID number of this BGP VRF.
BGP Instance	Autonomous System number and router ID.
NHT is Enabled	NHT Network enables the measurement and comparison of performance.
Recvd Msg count from RIB	Number of received change-messages from the RIB.
NHT delay-timer remaining seconds	Time remaining until the next decision cycle.
BGP nexthop(s)	Nexthop in the BGP to reach a certain destination.
Total number of IPV4 nexthops	Number of nexthops in the IPv4 Address Family.
Total number of IPV6 nexthops	Number of nexthops in the IPv6 Address Family.

show bgp nexthop-tree-details

Use this command to display BGP nexthop-tree details.

Command Syntax

```
show bgp nexthop-tree-details
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp nexthop-tree-details
BGP Instance: (Default), AS: 65534, router-id 51.1.1.3
AFI_IP Nexthop count : 0
AFI_IP6 Nexthop count : 0

BGP Instance: (Default), AS: 0, router-id 51.1.1.3
AFI_IP Nexthop count : 0
AFI_IP6 Nexthop count : 0
```

Table 12 explains the output fields.

Table 12: show bgp nexthop-tree-details output details

Field	Description
Bgp Instance	The Autonomous System number and router ID.
AFI_IP Nexthop count	Nexthop count for the IPv4 Address Family
AFI_IP6 Nexthop count	Nexthop count for the IPv6 Address Family

show bgp paths

Use this command to display BGP path information.

Command Syntax

```
show bgp paths
show bgp (ipv6) paths
show ip bgp paths
show ip bgp ipv4 (unicast|multicast) paths
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp paths

Address         Refcnt    Path
[0x1af8ee0:0]    (21)
[0x1b3ceb0:345] (14)      64602
[0x1c71d40:821] (12008)   64602 65500
[0x1d03fb0:822] (12008)   64602 65501
```

Table 13 explains the output fields.

Table 13: show bgp paths output details

Field	Description
Address	Hash and hash key separated by the colon character.
Refcnt	Number of routed using that path.
Path	Autonomous System Number (ASN) for the route.

show bgp prefix-list

Use this command to display routes matching the prefix-list.

Command Syntax

```
show bgp prefix-list WORD
show bgp (ipv4|ipv6) (unicast|multicast|) prefix-list WORD (vrf
  (VRFNAME|all|default|))
show ip bgp prefix-list WORD
show ip bgp prefix-list WORD (exact-match)
show ip bgp prefix-list WORD (exact-match) (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) prefix-list WORD
show ip bgp ipv4 (unicast|multicast) prefix-list WORD (exact-match)
```

Parameters

`WORD`	Name of the IP prefix list
`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`exact-match`	Exact match of the prefix list
`vrf`	VPN Routing/Forwarding instance
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp prefix-list Route1
BGP table version is 1, local router ID is 12.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop          Metric    LocPrf       Weight   Path
*>1.1.1.1/32   0.0.0.0            0         100         32768      i

Total number of prefixes 1
```

show bgp quote-regexp

Use this command to display route matching an AS path quoted regular expression.

Command Syntax

```
show bgp quote-regexp WORD
show bgp (ipv4|ipv6) (unicast|multicast|) quote-regexp WORD
show ip bgp quote-regexp WORD
show ip bgp ipv4 (unicast|multicast) quote-regexp WORD
```

Parameters

`WORD`	A regular expression to match the AS paths. Use quotes to enclose the regular expression.
`ipv4`	IPv4 route information
`ipv6`	IPv6 route information
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp quote-regexp ^$
BGP table version is 1, local router ID is 102.67.98.95
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network               Next Hop     Metric    LocPrf       Weight    Path
*> 2001:43f8:bb1::/64    :            0        100         32768      ?

Total number of prefixes 1
```

show bgp regexp

Use this command to display routes matching the AS path regular expression.

Command Syntax

```
show bgp regexp LINE
show bgp (ipv4|ipv6) (unicast|multicast) regexp LINE
show bgp (ipv6) regexp LINE
show ip bgp regexp LINE
show ip bgp vrf (VRFNAME|all|default) regexp LINE
show ip bgp ipv4 (unicast|multicast) regexp LINE
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`LINE`	A regular expression to match the AS paths
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp regexp ^$
BGP table version is 1, local router ID is 102.67.98.95
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network                Next Hop   Metric    LocPrf       Weight   Path
*> 2001:43f8:bb1::/64     ::        0        100          32768     ?

Total number of prefixes 1
```

show bgp route-map

Use this command to display routes that match the specified route map.

Command Syntax

```
show bgp route-map WORD
show bgp (ipv4|ipv6) (unicast|multicast) route-map WORD
show bgp (ipv6) route-map WORD
show ip bgp route-map WORD
show ip bgp route-map WORD (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) route-map WORD
```

Parameters

`WORD`	Routes matching the route-map
`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`vrf`	VPN routing/forwarding instance
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp prefix-list Route1
BGP table version is 1, local router ID is 12.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop            Metric    LocPrf       Weight  Path
*>1.1.1.1/32     0.0.0.0               0         100         32768    i
```

show bgp statistics

Use this command to display BGP statistics.

Command Syntax

```
show bgp statistics
```

Parameters

None

Command Mode

Privileged Exec and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp statistics

============================================================
BGP VRF default statistics
============================================================
Neighbor aggregated statistics (sent/received)
Msgs          Bytes           Opens          Updates
16/17           394/0           1/2           0/0
Keepalives  Notifications  Route-refresh  Capabilities
15/15            0/0           0/0           0/0
BGP I/O Information
Active Open attempts            : 0
Passive Open attempts           : 0
BGP I/O Open loops              : 0
BGP I/O Open calls              : 0
BGP I/O Open recv calls         : 0
BGP I/O Send calls              : 0
BGP I/O Recv calls              : 0
BGP I/O Write calls             : 0
BGP I/O Write loops             : 0
BGP I/O Write loop yields       : 0
BGP I/O Read calls              : 0
BGP I/O Read loops              : 0
BGP I/O Read loop yields        : 0
BGP I/O process nlri yields     : 0
BGP I/O process withdraw yields : 0
BGP Read time exceeded          : 0
BGP Update send pending         : 0
BGP Update buffer not available : 0
BGP Update walk suspended       : 0
BGP Yielded in updates          : 0
BGP Yielded in packing          : 0
BGP No sendbuf for peer         : 0
BGP No withdraw buf for peer    : 0
BGP Yields in update peer loop  : 0
```

```
No updates pending or no buffers: 0
No data to write                : 0
Msg queue recv errors           : 0
Sockets create/accept/close     : 2/1/2
Sockets create retries/failures : 1/0
Socket fd-close session         : 0
MemPool - Advertise             : | Total (0/0) blk_size:64
MemPool - AdjOut                : | Total (0/0) blk_size:12
MemPool - Advertise Attr        : | Total (0/0) blk_size:24
MemPool - BGP Info              : | Total (0/0) blk_size:216
MemPool - BGP Attr              : | Total (0/0) blk_size:224
MemPool - BGP Node IPv4         : | Total (0/0) blk_size:128
MemPool - BGP Node IPv6         : | Total (0/0) blk_size:136
MemPool - BGP Node EVPN         : | Total (0/0) blk_size:160
MemPool - BGP Node Max KeyLen   : | Total (0/0) blk_size:176
MemPool - BGP RIB msg4          : | Total (0/0) blk_size:4440
MemPool - BGP RIB msg6          : | Total (0/0) blk_size:424

#
```

show bgp summary

Use this command to display a summary of BGP neighbor status.

Command Syntax

```
show bgp summary
show bgp (ipv4|ipv6) (unicast|multicast|) summary
show ip bgp summary
show ip bgp summary (vrf (VRFNAME|all|default))
show ip bgp ipv4 (unicast|multicast) summary
```

Parameters

`ipv4`	IPv4 routes
`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp summary

BGP router identifier 6.6.6.6, local AS number 64601

BGP table version is 1

1 BGP AS-PATH entries

0 BGP community entries

Neighbor V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd

*12.1.1.24 64902    7          7       1      0      0   00:02:54       0

* Dynamically created based on a listen range command

BGP dynamic peer-group: group1
```

```
listen range:  12.1.0.0/16
Total number of dynamically created neighbors/limit: 1/(200)
Total number of dynamically created neighbors: 1
Total number of activated dynamic peer-groups for IPv4 Unicast address-family:
Total number of neighbors 1
Total number of Established sessions 1

BGP dynamic peer-group: group2
listen range:  12.2.0.0/16
Total number of dynamically created neighbors/limit: 0/(200)
Total number of dynamically created neighbors: 0
Total number of activated dynamic peer-groups for IPv4 Unicast address-family: 1
```

Header

```
BGP router identifier 10.10.15.50, local AS number 65000
1 BGP AS-PATH entries
0 BGP community entries
```

- The BGP router identifier is 10.10.15.50 and the local router AS number is 65000.
- The BGP table version tracks the local BGP table version. Any time the BGP best path algorithm executes, the table version increments.
- There is one BGP AS-PATH entry and no community entries.

Neighbor Entry Fields

Table 14 explains the fields for each neighbor entry.

Table 14: neighbor entry fields

Field	Description
Neighbor	IP address of peer.
V	BGP version of peer.
AS	Autonomous system number of peer.
MsgRcvd	Messages received since the BGP connection was established.
MsgSent	Messages sent since the BGP connection was established.
TblVer	Last version of the local router's BGP database advertised to the peer.
InQ	Received messages waiting in the input queue for further processing.
OutQ	Messages waiting in the output queue to be sent.

Table 14: neighbor entry fields (Continued)

Field	Description
Up/Down	Connection up time and down time.
State/PfxRcd	If the TCP session is up and the BGP peers have formed an adjacency, this field shows how many prefixes have been received from the remote neighbor. Other states: Idle: The local router has not allocated resources for the peer connection, so incoming connection requests are refused Idle (Admin): The peer has shut down Idle (PfxCt): Prefix overflow Idle (G-shut): Graceful shutdown Connect: BGP is waiting for the TCP connection to complete Active: the local router is trying to establish a TCP connection to the remote peer. You might see this if the local peer has been configured, but the remote peer is unreachable or has not been configured. OpenSent: BGP is waiting for an open message from its peer OpenConfirm: BGP received an open message from the peer and is now waiting for a keepalive or notification message. If BGP receives a keepalive message from the peer, the state changes to established. If the message is a notification, the state changes to idle. Established: BGP is ready to exchange update, notification, and keepalive messages with its peer Invalid: The session state is invalid

Neighbor Entry Example

```
10.10.14.51   4   100  93       120     0      0    0     00:42:16        0
```

- The neighbor has the IP address 10.10.14.51 and AS number 100.
- The neighbor uses BGP version 4.
- 93 messages have been received.
- 120 messages have been sent.
- The BGP routing table version is 0.
- There are no received messages waiting in the input queue for further processing.
- There are no messages waiting in the output queue to be sent.
- The connection has been up for 0 hours, 42 minutes and 53 seconds.
- The local router has received no prefixes from this neighbor.

show bgp X:X::X:X

Use this command to display BGP network information in an IPv6 environment.

Command Syntax

```
show bgp X:X::X:X
show bgp (ipv6) X:X::X:X
show bgp (ipv6) (unicast|multicast) X:X::X:X
```

Parameters

`ipv6`	IPv6 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`X:X::X:X`	IPv6 prefix (network), for example, 2003::

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bgp ipv6 6666::/64
BGP routing table entry for 6666::/64
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to peer-groups:
  Book-v6
  Local
    :: from :: (102.67.98.95)
      Origin incomplete,metric 0, localpref 100, weight 32768      valid,
sourced, best, source safi: 0
      Community: 37721:2000
      Last update: Thu Feb  4 09:00:50 2021
```

show bgp X:X::X:X/M longer prefixes

Use this command to display BGP network information along with mask information.

Command Syntax

```
show bgp X:X::X:X/M longer-prefixes
```

Parameters

`X:X::X:X/M`	IPv6 prefix (network/length), for example, 2003::/16

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show bgp 2001:43f8:bb1::/64 longer-prefixes
BGP table version is 1, local router ID is 102.67.98.95
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network               Next Hop   Metric    LocPrf       Weight  Path
*>2001:43f8:bb1::/64    ::         0         100         32768     ?

Total number of prefixes 1
```

show debugging bgp

Use this command to display BGP debugging options.

Command Syntax

```
show debugging bgp
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output from the show debugging bgp command.

```
#show debugging bgp
BGP debugging status:
 BGP debugging is on
 BGP events debugging is on
 BGP updates debugging is on
 BGP fsm debugging is on
```

show ip bgp

Use this command to display BGP routes.

Command Syntax

```
show ip bgp
show ip bgp ipv4 (unicast|multicast)(vrf (VRFNAME|all|default))
```

Parameters

`ipv4`	IPv4 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes
`VRFNAME`	VPN routing/forwarding instance name
`all`	All VRFs
`default`	Default VRF

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows routes learned from both iBGP and eBGP peers.

```
#show ip bgp
BGP table version is 0, local router ID is 10.100.0.77
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
   *> 172.16.1.0/24     10.10.10.78                          0       1 4 i
   *> 192.16.1.0        10.10.10.78       200                0       1 4 ?
   *                    10.100.0.62       100                0       3 4 ?
   *>i 192.17.1.0       10.100.0.62                 100      0       i

Total number of prefixes 2
```

Header

```
BGP table version is 0, local router ID is 10.100.0.77
```

- The BGP table version tracks the local BGP table version. Any time the BGP best path algorithm executes, the table version increments.
- The Router ID of the local router is `10.100.0.77`.

```
Status codes: s suppressed, d damped, h history, p stale, * valid, > best, i - internal
```

Table 15 explains the status codes in the header.

Table 15: status codes

Status code	Description	Comments
s	suppressed	Whether the route is suppressed and will not be advertised to the neighbors.
d	damped	When the penalty of a flapping route exceeds the suppress limit, the route is damped and remains in a withdrawn state until its penalty decreases below the reuse limit.
h	history	When the penalty of a flapping route does not exceed the suppress limit, the route is not damped and BGP maintains a history of the flapping route.
p	stale	When the BGP neighbor from which a route is learned is in graceful restart, the route is retained in the BGP routing table, but marked as stale.
*	valid	Whether the route is valid. When a route is not suppressed, damped, or present in the history, it is valid.
>	best	The route selected as the best path and installed in the kernel routing table.
i	internal	Whether the route is learned from an iBGP peer. If this symbol is not present, the route was learned from an eBGP peer.

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Table 16 explains the codes are at the end of each routing entry that show where the route originated.

Table 16: origin codes

Origin code	Description	Comments
i	IGP	The route originated from an Interior Gateway Protocol.
e	EGP	The route originated from an Exterior Gateway Protocol.
?	incomplete	Origin not known. Typically, these are routes redistributed from an Interior Gateway Protocol.

Route Entry Fields

Table 17 explains the fields shows for each route.

Table 17: route entry fields

Field	Description
Network	Network prefix installed in BGP. If multiple routes exist for the same prefix, only the first prefix is identified and others have blank spaces. The status codes are explained in Table 15.
Next Hop	IP address of the nexthop for this route.

Table 17: route entry fields (Continued)

Field	Description
Metric	Multiple-Exit Discriminator (MED). If there are multiple paths to the same destination from a single routing protocol, then the multiple paths have the same administrative distance and the best path is selected based on this metric. The path with the lowest metric is selected as the optimal path and installed in the routing table.
LocPrf	Local preference set with the `set local-preference` command. This value is used only with iBGP sessions within the local autonomous system to determine if a route towards a destination is the “best” one. The path with the highest local preference is preferred.
Weight	This field applies only to routes within an individual router. If a route was learned from a peer, it has a default weight of 0. All routes generated by the local router have a weight of 32,768.
Path and origin	The autonomous systems through which the prefix advertisement passed. The origin codes are explained in Table 16

Route Entry Examples

```
*> 172.16.1.0/24     10.10.10.78                            0    1 4 i
```

- The absence of status code “i” means the route is external and was learned from an eBGP peer.
- The “>” means this route is selected to be installed in the kernel routing table. Its network address is 172.16.1.0/24.
- The IP address of the nexthop for this route is 10.10.10.78.
- This route was learned from a peer, so it has a default weight of 0.
- The path “1 4” means the prefix advertisement passed through AS1 and AS4.
- The origin code “i” means the prefix was added by a network statement at an originating AS.

```
*> 192.16.1.0       10.10.10.78            200            0    1 4 ?
*                   10.100.0.62            100            0    3 4 ?
```

- The same prefix was learned from two different ASs, AS1 and AS3.
- The route learned from AS1 is chosen as the best route because AS1 has a lower Router ID (10.10.10.78) than AS2 (10.100.0.62). Although the metric of the route learned from AS1 is higher (200) than the route learned from AS3 (100), this attribute is not used in the best path selection decision because the metrics are compared only if the first (neighboring) AS is the same in the two paths.
- The origin code “?” indicates that the routes were learned through redistribution.

```
*>i192.17.1.0       10.100.0.62                    100     0    i
```

- The route is learned through an iBGP peer as indicated by the status code “i”.
- The preference of the route, used only with the local AS, is 100 (the default value).

show ip bgp cidr-only

Use this command to display routes with non-natural network masks.

Command Syntax

```
show ip bgp cidr-only
show ip bgp ipv4 (unicast|multicast) cidr-only
```

Parameters

`ipv4`	IPv4 routes
`unicast`	Unicast prefixes
`multicast`	Multicast prefixes

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output from the `show ip bgp cidr-only` command.

```
#show ip bgp cidr-only
BGP table version is 0, local router ID is 10.10.10.50
Status codes: s suppressed, d damped, h history, p stale, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 3.3.3.0/24       10.10.10.10                              0 11 i
Total number of prefixes 2
```

Table 18 explains the status codes in the header.

Table 18: status codes

Status code	Description	Comments
s	suppressed	Whether the route is suppressed and will not be advertised to the neighbors.
d	damped	When the penalty of a flapping route exceeds the suppress limit, the route is damped and remains in a withdrawn state until its penalty decreases below the reuse limit.
h	history	When the penalty of a flapping route does not exceed the suppress limit, the route is not damped and BGP maintains a history of the flapping route.
p	stale	When the BGP neighbor from which a route is learned is in graceful restart, the route is retained in the BGP routing table, but marked as stale.

Table 18: status codes (Continued)

*	valid	Whether the route is valid. When a route is not suppressed, damped, or present in the history, it is valid.
>	best	The route selected as the best path and installed in the kernel routing table.
i	internal	Whether the route is learned from an iBGP peer. If this symbol is not present, the route was learned from an eBGP peer.

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Table 19 explains the codes are at the end of each routing entry that show where the route originated.

Table 19: origin codes

Origin code	Description	Comments
i	IGP	The route originated from an Interior Gateway Protocol.
e	EGP	The route originated from an Exterior Gateway Protocol.
?	incomplete	Origin not known. Typically, these are routes redistributed from an Interior Gateway Protocol.

Route Entry Fields

Table 20 explains the fields shows for each route.

Table 20: route entry fields

Field	Description
Network	Network prefix installed in BGP. If multiple routes exist for the same prefix, only the first prefix is identified and others have blank spaces. The status codes are explained in Table 18.
Next Hop	IP address of the nexthop for this route.
Metric	Multiple-Exit Discriminator (MED). If there are multiple paths to the same destination from a single routing protocol, then the multiple paths have the same administrative distance and the best path is selected based on this metric. The path with the lowest metric is selected as the optimal path and installed in the routing table.
LocPrf	Local preference set with the `set local-preference` command. This value is used only with iBGP sessions within the local autonomous system to determine if a route towards a destination is the “best” one. The path with the highest local preference is preferred.
Weight	This field applies only to routes within an individual router. If a route was learned from a peer, it has a default weight of 0. All routes generated by the local router have a weight of 32,768.
Path and origin	The autonomous systems through which the prefix advertisement passed. The origin codes are explained in Table 19

show ip bgp community-info

Use this command to list all BGP community information.

Command Syntax

```
show ip bgp community-info
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp community-info
Address Refcnt Community
[0x262a4c0] (1) 1:1
```

show ip bgp peer-group

Use this command to list the BGP peer group information in ipv4 unicast.

Command Syntax

```
show ip bgp peer-group (WORD|)
```

Parameters

NAME	Name of the peer group.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp peer-group group1
BGP dynamic peer-group is group1, EBGP, remote AS 64902
  BGP dynamic peer-group group1 listen range group members:
  12.1.0.0/16
  BGP version 4
Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
  Peer-group member:
  *12.1.1.2
  Index 0, Offset 0, Mask 0x1
  0 accepted prefixes, 0 announced prefixes
```

show ip bgp peer-group vrf all

Use this command to list all BGP peer group VRF information.

Command Syntax

```
show ip bgp peer-group vrf all
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp peer-group vrf all
% VRF: VRF1

BGP dynamic peer-group is group2, EBGP, remote AS 64902
  BGP dynamic peer-group group2 listen range group members:
  12.2.0.0/16
  BGP version 4
Minimum time between advertisement runs is 30 seconds
% VRF: management
Peer-Group not found

% VRF: default

BGP dynamic peer-group is group1, EBGP, remote AS 64902
  BGP dynamic peer-group group1 listen range group members:
  12.1.0.0/16
  BGP version 4
Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
  Peer-group member:
  *12.1.1.2
  Index 0, Offset 0, Mask 0x1
  0 accepted prefixes, 0 announced prefixes
```

show ip bgp rtfilter all

Use this command to display route target filters sent and received.

Command Syntax

```
show ip bgp rtfilter all
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp rtfilter all
RTFilter's Received
********************
peer-ip 1.1.1.1
RTFilter's Sent
********************
peer-ip 1.1.1.1
```

show ip bgp scan

Use this command to display BGP scan status.

Command Syntax

```
show ip bgp scan
```

Parameters

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp scan
BGP VRF: (Default) VRF_ID 0
BGP scan interval is 60 secs
scan remain-time: 3 secs
Current BGP nexthop cache:
BGP connected route:
 10.10.10.0/24
 10.10.11.0/24
```

show ip extcommunity-list

Use this command to display BGP routes that match an extended community list.

Command Syntax

```
show ip extcommunity-list (WORD | )
show ip extcommunity-list (<1-199>|WORD)
show ip bgp extcommunity-list WORD (exact-match|)(vrf VRFNAME|)
```

Parameters

WORD	Name of extended community list
<1-199>	Number of extended community list
VRFNAME	VPN routing/forwarding instance name

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip extcommunity-list test
Named extended community expanded list test
    permit rt 100:100
```

show ip protocols

Use this command to display information about the IP protocols such as IP routing process parameters and statistics.

Command Syntax

```
show ip protocols
show ip protocols bgp
```

Parameters

bgp	BGP information

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip protocols bgp
Routing Protocol is "bgp 100"
Sending updates every 30 seconds with +/-50%, next due in 12 seconds
Timeout after 180 seconds, garbage collect after 120 seconds
Outgoing update filter list for all interface is not set
Incoming update filter list for all interface is not set
Default redistribution metric is 1
Redistributing: connected static
Default version control: send version 2, receive version 2
Interface        Send  Recv   Key-chain
   eth0                 2     2
Routing for Networks:
  10.10.0.0/24
Routing Information Sources:
  Gateway          BadPackets BadRoutes  Distance Last Update
 Distance: (default is 120)
```

Table 21 explains the fields shows for each route.

Table 21: show ip protocols output details

Field	Description
Routing Protocol is "bgp 100"	Specifies the routing protocol used.
Sending updates every 30 seconds	Specifies the time between sending updates.
Next due in 12 seconds	Precisely when the next update is due to be sent.
Timeout after 180 seconds	Specifies the value of the timeout parameter.
Redistributing	Lists the protocol that is being redistributed.

Table 21: show ip protocols output details

Field	Description
Routing for Networks	Specifies the networks for which the routing process is currently injecting routes.
Routing Information Sources	Lists all the routing sources the IP Infusion software is using to build its routing table.

show ip vrf

Use this command to display the routing information of the VRF, such as interface, route distinguisher, route-target, and so on.

Command Syntax

```
show ip vrf
show ip vrf WORD
```

Parameter

WORD	VRF name

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip vrf VRF_A
VRF VRF_A; (table=1)
```

show running-config as-path access-list

Use this command to show the running system status and configuration details for access lists based on autonomous system paths.

Command Syntax

```
show running-config as-path access-list
```

Parameters

None

Command Mode

Privileged exec mode, configure mode, router-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#show running-config as-path access-list
!
ip as-path access-list wer permit knsmk
!
(config)#
```

show running-config bgp

Use this command to show the running system configuration for BGP.

This command sorts address-families entries by VRF name in alphabetical order. The address-families type IPv4 have precedence to be displayed rather than address-families type IPv6 entries.

Command Syntax

```
show running-config bgp
```

Parameters

None

Command Mode

Privileged Exec mode, Configure mode, and Router mode.

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
OcNOS#show running-config bgp
!
router bgp 100
 neighbor 10.1.1.2 remote-as 100
 !
 address-family ipv4 vrf red
 exit-address-family
 !
 address-family ipv4 vrf yellow
 exit-address-family
 !
 address-family ipv6 vrf amber
 exit-address-family
 !
 address-family ipv6 vrf green
 exit-address-family
 !
 exit
!
OcNOS#
```

show running-config community-list

Use this command to show the running system status and configuration details for community lists.

Command Syntax

```
show running-config community-list
```

Parameters

None

Command Mode

Privileged exec mode, configure mode, router-map mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
(config)#show running-config community-list
!
ip community-list standard aspd permit internet
ip community-list expanded cspd deny ljj
ip community-list expanded cspd permit dcv
ip community-list expanded wde permit njhd
ip community-list expanded wer deny sde
(config)#
```

show running-config vrf

Use this command to show the running system configuration details for VRF instances.

This command sorts VRF entries by name in alphabetical order.

Command Syntax

```
show running-config vrf
```

Parameters

None

Command Mode

Privileged Exec mode, Configure mode, and Router mode.

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
OcNOS#show running-config vrf
ip vrf management
!
ip vrf test2
!
ip vrf test5
!
mac vrf test1
!
mac vrf test3
!
mac vrf test4
!
!
end

OcNOS#
```

show running-config ip vrf

Use this command to show the running system status and configuration details for IP VRF entries.

This command sorts VRF entries by name in alphabetical order.

Command Syntax

```
show running-config ip vrf
```

Parameters

None

Command Mode

Privileged Exec mode, Configure mode, and Router mode.

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
OcNOS#show running-config ip vrf
ip vrf management
!
ip vrf test2
!
ip vrf test5
!
!
end
OcNOS#
```

show running-config mac vrf

Use this command to show the running system status and configuration details for VRF MAC entries.

This command sorts VRF entries by name in alphabetical order.

Command Syntax

```
show running-config mac vrf
```

Parameters

None

Command Mode

Privileged Exec mode, Configure mode, and Router mode.

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
OcNOS#show running-config mac vrf
mac vrf test1
!
mac vrf test3
!
mac vrf test4
!
!
end
OcNOS#
```

show running-config vrf word

Use this command to show the running system status and configuration details for a specified VRF instance name.

This command displays the VRF instance association to interfaces and protocols.

Command Syntax

```
show running-config vrf word
```

Parameters

word	Virtual Routing and Forwarding name

Command Mode

Privileged Exec mode, Configure mode and Router mode.

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

```
OcNOS#show running-config vrf test2
ip vrf test2
!
interface eth2
 ip vrf forwarding test2
!
router bgp 100
 address-family ipv4 vrf test2
 exit-address-family
!
end

OcNOS#
```

Appendix D Regular Expressions

Table 4-1 shows the regular expression special characters used in BGP commands. You can use these characters in combination to build complex regular expressions.

Table 4-1: Regular expression characters

Symbol	Character	Meaning
^	Caret	Matches the beginning of the input string. When used at the beginning of a string of characters, it negates a pattern match.
$	Dollar sign	Matches the end of the input string.
.	Period	Matches a single character (including white spaces).
*	Asterisk	Matches none or more sequences of a pattern.
+	Plus sign	Matches one or more sequences of a pattern.
?	Question mark	Matches none or one occurrence of a pattern.
_	Underscore	Matches spaces, commas, braces, parenthesis, or the beginning and end of an input string.
[]	Brackets	A range of single-characters.
-	Hyphen	Separates the end points of a range.

Bidirectional Forwarding Detection Configuration

CHAPTER 1 Base BFD Configuration

This chapter provides the steps for configuring the base Bidirectional Forwarding Detection (BFD) setup.

Overview

This section provides an overview of Bidirectional Forwarding Detection (BFD). BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols like BGP, EIGRP, IS-IS, and OSPF. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators.The configuration and command reference for BFD is given in the following chapters in this section.

Note: It is mandatory to enable the bfd-group filter for bfd session used in Trident-III devices. This filter is applicable only for Trident-III devices. Refer to hardware-profile filter (XGS) CLI section for more information.

Topology

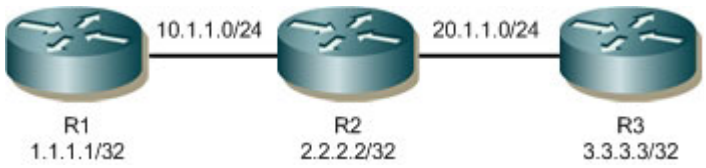

Figure P-6: Basic Topology of Three Routers

BFD Echo Function

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#bfd session 10.1.1.1 10.1.1.2`	Configure BFD session on interface.
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#bfd echo`	Enable BFD echo mode.
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit from interface mode.

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#bfd session 10.1.1.2 10.1.1.1`	Configure BFD session on interface.
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#bfd echo`	Enable BFD echo mode.

`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit from interface mode.

Validation

R1

```
#sh bfd session detail

Session Interface Index : 3                 Session Index : 1
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                   Session State : Down
Local Discriminator : 1                     Local Address : 10.1.1.1/32
Remote Discriminator : 1                    Remote Address : 10.1.1.2/32
Local Port : 49152                          Remote Port : 3784
Local Echo Port : 49153
Options :
Echo Enabled
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 2000              Neg Rx: 2000              Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 1000
Storage type : 2
Sess down time : 00:00:01
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 000000000000010a                   Pkt Out : 00000000000001f7
Echo Out : 00000000000000ab                 IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000              IPv6 Pkt Out : 0000000000000000
UP Count : 57                               UPTIME : 00:00:00

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

R2

```
#sh bfd session detail

Session Interface Index : 3                 Session Index : 1
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                   Session State : Up
Local Discriminator : 1                     Local Address : 10.1.1.2/32
Remote Discriminator : 1                    Remote Address : 10.1.1.1/32
Local Port : 49152                          Remote Port : 3784
Local Echo Port : 49153
Options :
Echo Enabled
Diagnostics : None
```

```
Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 2000              Neg Rx: 2000              Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 1000
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000ad                Pkt Out : 0000000000000105
Echo Out : 0000000000000063              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 33                            UPTIME : 00:00:03

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

BFD Slow Timer

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#bfd session 10.1.1.1 10.1.1.2`	Configure BFD session on interface.
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#bfd slow-timer 1000`	Configure BFD `slow-timer` in milliseconds.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit from interface mode.

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#bfd session 10.1.1.2 10.1.1.1`	Configure BFD session on interface.
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#bfd slow-timer 1000`	Configure BFD `slow-timer` in milliseconds.

BFD Multihop Peer Timer

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#router ospf 1`	Enter router OSPF mode
`R1(config-router)#network 10.1.1.0/24 area 0`	Advertise the network to area 0
`R1(config-router)#redistribute connected`	
`R1(config-router)#exit`	Exit router OSPF mode
`R1(config)#bfd multihop-peer 20.1.1.3 interval 100 minrx 100 multiplier 3`	Configure BFD `multihop-peer` timer and reception intervals in milliseconds and the Hello multiplier.
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit router OSPF mode

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#interface eth2`	Enter interface Mode.
`R2(config-if)#ip address 20.1.1.2/24`	Assign IP address for interface
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#router ospf 1`	Enter router OSPF mode
`R2(config-router)#network 10.1.1.0/24 area 0`	Advertise the network to area 0
`R2(config-router)#network 20.1.1.0/24 area 0`	Advertise the network to area 0
`R2(config-router)#redistribute connected`	
`R2(config-router)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-router)#exit`	Exit router OSPF mode

R3

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth2`	Enter interface Mode.
`R1(config-if)#ip address 20.1.1.3/24`	Assign IP address for interface
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#router ospf 1`	Enter router OSPF mode
`R1(config-router)#network 20.1.1.0/24 area 0`	Advertise the network to area 0
`R1(config-router)#redistribute connected`	
`R1(config-router)#exit`	Exit router OSPF mode

`R1(config)#bfd multihop-peer 10.1.1.1 interval 100 minrx 100 multiplier 3`	Configure BFD `multihop-peer` timer and reception intervals in milliseconds and the Hello multiplier.v
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit router OSPF mode

Validation:

R1

```
#sh bfd session detail


Session Interface Index : 0              Session Index : 2
Lower Layer : IPv4                       Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 2                  Local Address : 10.1.1.1/32
Remote Discriminator : 0                 Remote Address : 20.1.1.3/32
Local Port : 49153                       Remote Port : 4784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 100              Min Rx: 100              Multiplier: 3
Neg Tx: 0                Neg Rx: 0                Neg detect mult: 0
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 0000000000000000                Pkt Out : 0000000000000097
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 0                             UPTIME : 00:00:00

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

R3

```
#sh bfd session detail


Session Interface Index : 0              Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 1                  Local Address : 20.1.1.3/32
Remote Discriminator : 1                 Remote Address : 10.1.1.1/32
Local Port : 49152                       Remote Port : 4784
```

```
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 100               Min Rx: 100               Multiplier: 3
Neg Tx: 100               Neg Rx: 100               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000001df                Pkt Out : 00000000000001e0
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:01:26

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

BFD Single-hop Session Timer

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#bfd session 10.1.1.1 10.1.1.2`	Configure BFD session on interface.
`R1(config-if)#bfd interval 100 minrx 100 multiplier 4`	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit from interface mode.

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#bfd session 10.1.1.2 10.1.1.1`	Configure BFD session on interface.
`R2(config-if)#bfd interval 100 minrx 100 multiplier 4`	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exit from interface mode.

Validation:

R1

```
#sh bfd session detail

Session Interface Index : 3             Session Index : 1
Lower Layer : IPv4                      Version : 1
Session Type : Single Hop               Session State : Up
Local Discriminator : 1                 Local Address : 10.1.1.1/32
Remote Discriminator : 1                Remote Address : 10.1.1.2/32
Local Port : 49152                      Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 100               Min Rx: 100               Multiplier: 4
Neg Tx: 100               Neg Rx: 100               Neg detect mult: 4
Min echo Tx: 100          Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 000000000000156e               Pkt Out : 0000000000001748
Echo Out : 000000000000019b             IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000          IPv6 Pkt Out : 0000000000000000
UP Count : 139                          UPTIME : 00:14:23

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

R2

```
#sh bfd session detail

Session Interface Index : 3             Session Index : 1
Lower Layer : IPv4                      Version : 1
Session Type : Single Hop               Session State : Up
Local Discriminator : 1                 Local Address : 10.1.1.2/32
Remote Discriminator : 1                Remote Address : 10.1.1.1/32
Local Port : 49152                      Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 100               Min Rx: 100               Multiplier: 4
Neg Tx: 100               Neg Rx: 100               Neg detect mult: 4
Min echo Tx: 100          Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
```

```
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 000000000000181d                  Pkt Out : 00000000000019ab
Echo Out : 00000000000001b5                IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000             IPv6 Pkt Out : 0000000000000000
UP Count : 145                             UPTIME : 00:15:19

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

BFD Echo Interval

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#bfd session 10.1.1.1 10.1.1.2`	Configure BFD session on interface.
`R1(config-if)#bfd echo interval 100`	
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#bfd echo`	Enable BFD echo mode.
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit from interface mode.

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#bfd session 10.1.1.2 10.1.1.1`	Configure BFD session on interface.
`R1(config-if)#bfd echo interval 100`	.
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#bfd echo`	Enable BFD echo mode.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit from interface mode.

Validation

R1

```
#sh bfd session detail

Session Interface Index : 3             Session Index : 1
Lower Layer : IPv4                      Version : 1
Session Type : Single Hop               Session State : Down
Local Discriminator : 1                 Local Address : 10.1.1.1/32
Remote Discriminator : 1                Remote Address : 10.1.1.2/32
Local Port : 49152                      Remote Port : 3784
Local Echo Port : 49153
Options :
Echo Enabled
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 2000              Neg Rx: 2000              Neg detect mult: 3
Min echo Tx: 100          Min echo Rx: 1000         Neg echo intrvl: 1000
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000001ea               Pkt Out : 0000000000000398
Echo Out : 0000000000000147             IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000          IPv6 Pkt Out : 0000000000000000
UP Count : 109                          UPTIME : 00:00:00

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

R2

```
#sh bfd session detail

Session Interface Index : 3             Session Index : 1
Lower Layer : IPv4                      Version : 1
Session Type : Single Hop               Session State : Down
Local Discriminator : 1                 Local Address : 10.1.1.2/32
Remote Discriminator : 1                Remote Address : 10.1.1.1/32
Local Port : 49152                      Remote Port : 3784
Local Echo Port : 49153
Options :
Echo Enabled
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 2000              Neg Rx: 2000              Neg detect mult: 3
```

```
Min echo Tx: 100          Min echo Rx: 1000         Neg echo intrvl: 1000
Storage type : 2
Sess down time : 00:00:01
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 000000000000028f                Pkt Out : 00000000000003b8
Echo Out : 0000000000000183              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 129                           UPTIME : 00:00:00

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

CHAPTER 2 BFD Protocol Configurations

This chapter describes the BFD protocol configurations.

OSPF—BFD Single-Hop Session

This section provides the steps for configuring BFD for Single-Hop OSPF.

Topology

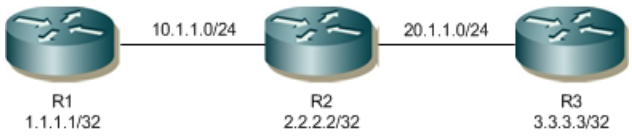

Figure P-7: Single-Hop OSPF Topology

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface lo`	Enter the Interface configuration mode for lo
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Assign IP Address
`R1(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP Address
`R1(config-if)#exit`	Exit the Interface configuration mode.
`R1(config)#router ospf 100`	Enter the Router mode for OSPF.
`R1(config-router)#network 10.1.1.0/24 area 1`	Advertise network `10.1.1.0/24` in OSPF `area 1`.
`R1(config-router)#network 1.1.1.1/32 area 1`	Advertise loopback network `1.1.1.1/32` in OSPF `area 1`.
`R1(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R1(config-router)#commit`	Commit the configuration on the node
`R1(config-router)#exit`	Exit the router mode

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode.
`R2(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R2(config-if)#ip address 20.1.1.1/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode
`R2(config)#router ospf 100`	Enter the Router mode for OSPF.

`R2(config-router)#network 10.1.1.0/24 area 1`	Advertise network `10.1.1.0/24` in OSPF `area 1`.
`R2(config-router)#network 20.1.1.0/24 area 1`	Advertise network `20.1.1.0/24` in OSPF `area 1`.
`R2(config-router)#network 2.2.2.2/32 area 1`	Advertise loopback network `2.2.2.2/32` in OSPF `area 1`.
`R2(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R2(config)#interface lo`	Enter the Interface configuration mode for lo
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Assign IP Address
`R2(config-router)#commit`	Commit the configuration on the node
`R2(config-router)#exit`	Exit the router mode.

R3

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#interface lo`	Enter the Interface configuration mode for lo
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Assign IP Address
`R3(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R3(config-if)#ip address 20.1.1.2/24`	Assign IP Address.
`R3(config-if)#exit`	Exit the Interface configuration mode.
`R3(config)#router ospf 100`	Enter the Router mode for OSPF.
`R3(config-router)#network 20.1.1.0/24 area 1`	Advertise network `20.1.1.0/24` in OSPF `area 1.`
`R3(config-router)#network 3.3.3.3/32 area 1`	Advertise loopback network `3.3.3.3/32` in OSPF `area 1.`
`R3(config-router)#bfd all-interfaces`	Enable BFD for all neighbors
`R3(config-router)#commit`	Commit the configuration on the node
`R3(config-router)#exit`	Exit the router mode.

Validation

R1

```
R1#show ip ospf n

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface   Instance ID
2.2.2.2           1   Full/Backup      00:00:35    10.1.1.2            xe23          0


R1#show bfd session

BFD process for VRF: (DEFAULT VRF)
========================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface  Down-
Reason  Remote-Addr
```

```
1                     1              IPv4        Single-Hop      Up         00:02:55
xe23       NA                  10.1.1.2/32
257                   0              IPv4        Multi-Hop       Up          00:00:18
xe8         NA                   3.3.3.3/32
Number of Sessions:    2

R1#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
====

Session Interface Index : 10023          Interface name :xe23
Session Index : 1
Lower Layer : IPv4                                Version : 1
Session Type : Single Hop                    Session State : Up
Local Discriminator : 1                         Local Address : 10.1.1.1/32
Remote Discriminator : 1                     Remote Address : 10.1.1.2/32
Local Port : 49152                                Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                    Multiplier: 3
Neg Tx: 250                    Neg Rx: 250                    Neg detect mult: 3
Min echo Tx: 1000          Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000814                Pkt Out : 00000000000000000814
Pkts Drop : 00000000000000000000          Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                 UPTIME : 00:02:59

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 0                   Session Index : 257
Lower Layer : IPv4                                Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 257                      Local Address : 1.1.1.1/32
Remote Discriminator : 0                     Remote Address : 3.3.3.3/32
Local Port : 49153                                 Remote Port : 4784
Options :
```

```
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                   Min Rx: 250                   Multiplier: 3
Neg Tx: 0                       Neg Rx: 0                       Neg detect mult: 0
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                Pkt Out : 00000000000000000107
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 0                                                UPTIME : 00:00:00
R3#show bfd session
BFD process for VRF: (DEFAULT VRF)
================================================================================================
Sess-Idx Remote-Disc Lower-Layer Sess-Type Sess-State UP-Time Interface    Down-Reason
Remote-Addr
1          1      IPv4       Single-Hop   Up     00:04:33   eth2           NA
20.1.1.3/32
------------------------------------------------------------------------------------------------

R1#show bfd session detail
=======================================================

Session Interface Index : 3
Interface name : eth1                    Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                  Local Address : 10.1.1.1/32
Remote Discriminator : 1                 Remote Address : 10.1.1.2/32
Local Port : 49152                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled
```

```
Counters values:
Pkt In : 0000000000000cf3                Pkt Out : 0000000000000cfc
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:12:12

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------
Number of Sessions:    1

R2#show bfd session detail
======================================================

Session Interface Index : 3
Interface name : eth1                    Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                  Local Address : 10.1.1.2/32
Remote Discriminator : 1                 Remote Address : 10.1.1.1/32
Local Port : 49152                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 2000             Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 0000000000000d6f                Pkt Out : 0000000000000da5
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:12:39

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------

Session Interface Index : 4
Interface name : eth1                    Session Index : 2
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                  Local Address : 20.1.1.1/32
Remote Discriminator : 1                 Remote Address : 20.1.1.2/32
```

```
Local Port : 49153                        Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 2000              Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 0000000000000d30                 Pkt Out : 0000000000000d5e
Echo Out : 0000000000000000               IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000            IPv6 Pkt Out : 0000000000000000
UP Count : 1                              UPTIME : 00:12:24
```

R2

```
R2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface  Instance ID
1.1.1.1           1   Full/DR          00:00:32      10.1.1.1           xe9/3       0
3.3.3.3           1   Full/Backup    00:00:40     20.1.1.2           xe17/1     0


R2#show bfd session

BFD process for VRF: (DEFAULT VRF)
=====================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface   Down-
Reason  Remote-Addr
1               1               IPv4            Single-Hop    Up         00:02:33
xe9/3           NA          10.1.1.1/32
2               1               IPv4            Single-Hop    Up         00:01:54
xe17/1          NA          20.1.1.2/32
Number of Sessions:    2


R2#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
=====================================================================================

Session Interface Index : 10023          Interface name :xe9/3
Session Index : 1
```

```
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                     Local Address : 10.1.1.2/32
Remote Discriminator : 1                 Remote Address : 10.1.1.1/32
Local Port : 49152                          Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                  Min Rx: 250                  Multiplier: 3
Neg Tx: 250                  Neg Rx: 250                  Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000713               Pkt Out : 00000000000000000714
Pkts Drop : 00000000000000000000          Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                UPTIME : 00:02:37

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 10053          Interface name :xe17/1
Session Index : 2
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                     Local Address : 20.1.1.1/32
Remote Discriminator : 1                 Remote Address : 20.1.1.2/32
Local Port : 49153                           Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                  Min Rx: 250                   Multiplier: 3
Neg Tx: 250                  Neg Rx: 250                  Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled
```

```
Counters values:
Pkt In : 00000000000000000535                Pkt Out : 00000000000000000537
Pkts Drop : 00000000000000000000             Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000              IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                 UPTIME : 00:01:58

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

R3

```
R3#show ip ospf n

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface   Instance ID
2.2.2.2           1     Full/DR         00:00:36    20.1.1.1          xe17/10       0


R3#show bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface   Down-
Reason  Remote-Addr
1                    2          IPv4            Single-Hop      Up         00:01:21
xe17/1             NA           20.1.1.1/32
257                    1           IPv4              Multi-Hop              Up
00:00:00      xe6                NA           1.1.1.1/32
Number of Sessions:    2
Total number of Established sessions 2


R3#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================

Session Interface Index : 10053          Interface name :xe17/1
Session Index : 1
Lower Layer : IPv4                            Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                     Local Address : 20.1.1.2/32
Remote Discriminator : 2                 Remote Address : 20.1.1.1/32
Local Port : 49152                            Remote Port : 3784
Options :

Diagnostics : None
```

```
Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                    Multiplier: 3
Neg Tx: 250                    Neg Rx: 250                    Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000391                Pkt Out : 00000000000000000391
Pkts Drop : 00000000000000000000          Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                UPTIME : 00:01:26

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 0                  Session Index : 257
Lower Layer : IPv4                                Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 257                     Local Address : 3.3.3.3/32
Remote Discriminator : 1                     Remote Address : 1.1.1.1/32
Local Port : 49153                                Remote Port : 4784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                        Min Rx: 250                    Multiplier: 3
Neg Tx: 250                        Neg Rx: 250                    Neg detect mult:250
Min echo Tx: 1000             Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                Pkt Out : 00000000000000000047
Pkts Drop : 00000000000000000000          Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 0                                                UPTIME : 00:00:00

Protocol Client Info:
```

```
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

OSPF—BFD Multi-Hop Session

This section provides the steps for configuring BFD for OSPF multi-hop sessions.

Topology

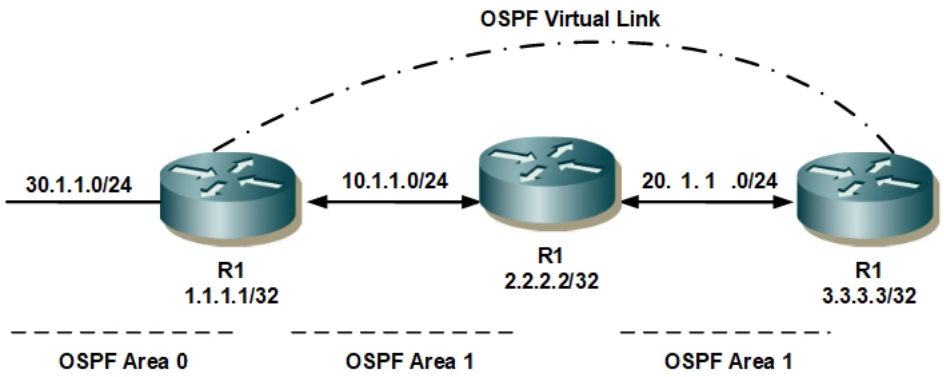

Figure P-8: Multi-hop OSPFv2 Topology

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface lo`	Enter the Interface configuration mode for lo
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Assign IP Address
`R1(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP Address
`R1(config-if)#exit`	Exit the Interface configuration mode.
`R1(config)#interface eth2`	Enter the Interface configuration mode for `eth1`.
`R1(config-if)#ip address 30.1.1.1/24`	Assign IP Address
`R1(config-if)#exit`	Exit the Interface configuration mode.
`R1(config)#router ospf 100`	Enter the Router mode for OSPF.
`R1(config-router)#ospf router-id 1.1.1.1`	OSPF router ID in IPv4 format
`R1(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R1(config-router)#network 10.1.1.0/24 area 1`	Advertise network `20.1.1.0/24` in OSPF `area 1`.
`R1(config-router)#network 1.1.1.1/32 area 1`	Advertise loopback network `1.1.1.1/32` in OSPF `area 1`.
`R1(config-router)#network 30.1.1.0/24 area 0`	Advertise network `30.1.1.0/24` in OSPF `area 0`.

`R1(config-router)#area 1 virtual-link 3.3.3.3 fall-over bfd`	Create a virtual link to R3 with BFD.
`R1(config-router)#commit`	Commit the configuration on the node

R2

`R2#configure terminal`	Enter the Configure mode
`R2(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode.
`R2(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R2(config-if)#ip address 20.1.1.1/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode
`R2(config)#interface lo`	Enter the Interface configuration mode for lo
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode
`R2(config)#router ospf 100`	Enter the Router mode for OSPF.
`R2(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R2(config-router)#ospf router-id 2.2.2.2`	OSPF router ID in IPv4 format
`R2(config-router)#network 10.1.1.0/24 area 1`	Advertise network `10.1.1.0/24` in OSPF `area 1`.
`R2(config-router)#network 20.1.1.0/24 area 1`	Advertise network `20.1.1.0/24` in OSPF `area 1`.
`R2(config-router)#network 2.2.2.2/32 area 1`	Advertise loopback network `2.2.2.2/32` in OSPF `area 1`.
`R2(config-router)#commit`	Commit the configuration on the node
`R2(config-router)#exit`	Exit the router mode.

R3

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#interface lo`	Enter the Interface configuration mode for lo
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Assign IP Address
`R3(config-if)#exit`	Exit the Interface configuration mode.
`R3(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R3(config-if)#ip address 20.1.1.2/24`	Assign IP Address.
`R3(config-if)#exit`	Exit the Interface configuration mode.
`R3(config)#router ospf 100`	Enter the Router mode for OSPF.
`R3(config-router)#ospf router-id 3.3.3.3`	OSPF router ID in IPv4 format
`R3(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R3(config-router)#network 20.1.1.0/24 area 1`	Advertise network `20.1.1.0/24` in OSPF `area 1`.
`R3(config-router)#network 3.3.3.3/32 area 1`	Advertise loopback network `3.3.3.3/32` in OSPF `area 1`.

© 2025 IP Infusion Inc. Proprietary

`R3(config-router)#area 1 virtual-link 1.1.1.1 fall-over bfd`	Create a virtual link to R1 with BFD.
`R3(config-router)#commit`	Commit the configuration on the node

Validation

```
R1#sh bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface
Down-Reason  Remote-Addr
257           257              IPv4       Multi-Hop       Up        00:00:54
NA         NA                   20.1.1.2/32

Number of Sessions:    1


R2#sh bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface
Down-Reason  Remote-Addr
Number of Sessions:    0



R3#sh bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface
Down-Reason  Remote-Addr
257       257         IPv4        Multi-Hop   Up         00:01:46  NA       NA
10.1.1.1/32

Number of Sessions:    1


R1#sh bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================

Session Interface Index : 0              Session Index : 257
Lower Layer : IPv4                       Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 257                Local Address : 10.1.1.1/32
Remote Discriminator : 257               Remote Address : 20.1.1.2/32
Local Port : 49152                       Remote Port : 4784
```

```
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000606              Pkt Out : 00000000000000000605
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 1                               UPTIME : 00:02:13

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------------
Number of Sessions:    1



R2#sh bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
Number of Sessions:    0
R2#


R3#sh bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================

Session Interface Index : 0             Session Index : 257
Lower Layer : IPv4                      Version : 1
Session Type : Multihop Arbit Path      Session State : Up
Local Discriminator : 257               Local Address : 20.1.1.2/32
Remote Discriminator : 257              Remote Address : 10.1.1.1/32
Local Port : 49152                      Remote Port : 4784
Options :

Diagnostics : None
```

```
Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000529              Pkt Out : 00000000000000000530
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 1                               UPTIME : 00:01:56

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

BFD Configuration in IS-IS

This section provides the steps for configuring BFD for the IS-IS protocol.

Topology

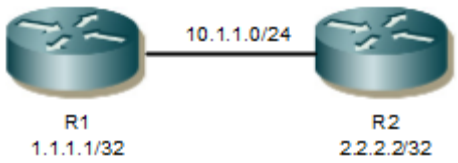

Figure P-9: Basic Topology for BFD-ISIS

R1

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth1`	Enter interface mode.
`R1(config-if)#ip address 10.1.1.1/24`	Configure IP address.
`R2(config-if)#ip router isis 1`	Configure ISIS router
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router isis 1`	Enter the Router mode for IS-IS.
`R1(config-router)#net 10.0000.0000.0001.00`	Advertise network `10.0000.0000.0001.00` in IS-IS.

`R1(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R1(config-router)#commit`	Commit the configuration on the node

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter interface mode.
`R2(config-if)#ip address 10.1.1.2/24`	Configure IP address.
`R2(config-if)#ip router isis 1`	Configure ISIS router
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#router isis 1`	Enter the Router mode for IS-IS.
`R2(config-router)#net 10.0000.0000.0002.00`	Advertise network `10.0000.0000.0002.00` in IS-IS.
`R2(config-router)#bfd all-interfaces`	Enable BFD for all neighbors.
`R2(config-router)#commit`	Commit the configuration on the node

Validation

```
R1#show bfd session

    BFD process for VRF: (DEFAULT VRF)
    ========================================================================
    =======
    Sess-Idx Remote-Disc Lower-Layer Sess-Type Sess-State UP-Time Interface
    Down-Reason Remote-Addr
    1                    1           IPv4             Single-Hop       Up
    00:00:42      eth1           NA                     10.1.1.1/32

    Number of Sessions:    1

    R2#show bfd session

    BFD process for VRF: (DEFAULT VRF)
    ========================================================================
    =======
    Sess-Idx Remote-Disc Lower-Layer Sess-Type Sess-State UP-Time Interface
    Down-Reason  Remote-Addr
    1            1            IPv4    Single-Hop  Up  00:10:23      eth1
    NA       10.1.1.2/32
    1            1            IPv4    Single-Hop  Up  00:10:23      eth2
    NA       20.1.1.2/32
    Number of Sessions:    2

R1#show bfd session detail
=======================================================

Session Interface Index : 3
Interface name : eth1                    Session Index : 2
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                  Local Address : 10.1.1.1/32
Remote Discriminator : 3                 Remote Address : 10.1.1.2/32
```

```
Local Port : 49153                      Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250             Min Rx: 250             Multiplier: 3
Neg Tx: 250             Neg Rx: 2000            Neg detect mult: 3
Min echo Tx: 1000       Min echo Rx: 1000       Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 000000000000027b               Pkt Out : 0000000000000278
Echo Out : 0000000000000000             IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000          IPv6 Pkt Out : 0000000000000000
UP Count : 1                            UPTIME : 00:02:19

Protocol Client Info:
ISIS-> Client ID: 6      Flags: 4
------------------------------------------------------------
Number of Sessions:    1

R2#sh bfd session detail
============================================================

Session Interface Index : 3
Interface name : eth1                   Session Index : 3
Lower Layer : IPv4                      Version : 1
Session Type : Single Hop               Session State : Up
Local Discriminator : 3                 Local Address : 10.1.1.2/32
Remote Discriminator : 2                Remote Address : 10.1.1.1/32
Local Port : 49154                      Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250             Min Rx: 250             Multiplier: 3
Neg Tx: 250             Neg Rx: 250             Neg detect mult: 3
Min echo Tx: 1000       Min echo Rx: 1000       Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000002db               Pkt Out : 00000000000002dd
```

```
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:02:42

Protocol Client Info:
ISIS-> Client ID: 6      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

BFD Configuration in BGP

This section provides the steps for configuring BFD for the BGP protocol.

Topology

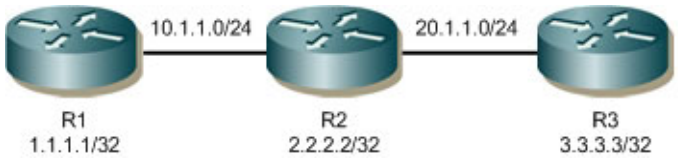

Figure P-10: Basic Topology for BFD in BGP

R1

`R1#configure terminal`	Enter configure mode
`R1(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP Address
`R1(config-if)#exit`	Exit the Interface configuration mode.
`R1(config)# interface lo`	Enter interface mode
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Configure the IP address on loopback interface.
`R1(config-if)#bfd session  1.1.1.1  3.3.3.3 multihop`	Enable BFD multihop session
`R1(config)#router ospf 100`	Enter Router mode for OSPF
`R1(config-router)# ospf router-id 1.1.1.1`	Configure router-id in OSPF
`R1(config-router)# network 1.1.1.1/32 area 0.0.0.1`	Advertise network 1.1.1.1/32 in OSPF area 1.
`R1(config-router)# network 10.1.1.0/24 area 0.0.0.1`	Advertise network 10.1.1.0/24 in OSPF area 1.
`R1(config)#router bgp 100`	Enter Router mode for BGP
`R1(config-router)# neighbor 3.3.3.3 remote-as 100`	Add the neighbor 3.3.3.3 to remote-as 100.
`R1(config-router)# neighbor 3.3.3.3 fall-over bfd`	Enable BFD option for neighbor `multihop`
`R1(config-router)#neighbor 3.3.3.3 update-source lo`	Add neighbor 3.3.3.3 to update-source lo.

`R1(config-router)#neighbor 10.1.1.2 remote-as 100`	Add neighbor 10.1.1.2 to remote-as 100.
`R1(config-router)# neighbor 10.1.1.2 fall-over bfd`	Enable BFD option for neighbor
`R1(config-router)#commit`	Commit the configuration on the node
`R1(config-router)#end`	Exit from Router BGP mode

R2

`R2#configure terminal`	Enter configure mode
`R2(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode.
`R2(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R2(config-if)#ip address 20.1.1.1/24`	Assign IP Address
`R2(config-if)#exit`	Exit the Interface configuration mode
`R2(config)# interface lo`	Enter interface mode
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Configure the IP address on loopback interface.
`R2(config-if)#exit`	Exit the Interface configuration mode
`R2(config)#router ospf 100`	Enter Router mode for OSPF
`R2(config-router)#ospf router-id 2.2.2.2`	Configure router-id in OSPF
`R2(config-router)# network 2.2.2.2/32 area 0.0.0.1`	Advertise network 1.1.1.1/32 in OSPF area 1.
`R2(config-router)#network 10.1.1.0/24 area 0.0.0.1`	Advertise network 10.1.1.0/24 in OSPF area 1.
`R2(config-router)# network 20.1.1.0/24 area 0.0.0.1`	Advertise network 20.1.1.0/24 in OSPF area 1.
`R3(config-router)#exit`	Exit from Router OSPF mode
`R2(config)#router bgp 100`	Enter Router mode for BGP
`R2(config-router)# neighbor 10.1.1.1 remote-as 100`	Add neighbor 10.1.1.1 to remote-as 100.
`R2(config-router)#neighbor 10.1.1.1 fall-over bfd`	Enable BFD option for neighbor
`R2(config-router)# neighbor 20.1.1.2 remote-as 100`	Add neighbor 20.1.1.2 to remote-as 100.
`R2(config-router)# neighbor 20.1.1.2 fall-over bfd`	Enable BFD option for neighbor
`R2(config-router)#commit`	Commit the configuration on the node
`R2(config-router)#end`	Exit from Router BGP mode

R3

`R3#configure terminal`	Enter configure mode
`R3(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.

R3(config-if)#ip address 20.1.1.2/24	Assign IP Address.
R3(config-if)#exit	Exit the Interface configuration mode.
R3(config)# interface lo	Enter interface mode
R3(config-if)#ip address 3.3.3.3/32 secondary	Configure IP address on loopback interface.
R1(config-if)#bfd session 3.3.3.3 1.1.1.1 multihop	Enable BFD multihop session
R2(config-if)#exit	Exit the Interface configuration mode
R3(config)#router ospf 100	Enter Router mode for OSPF
R3(config-router)#ospf router-id 3.3.3.3	Configure router-id in OSPF
R3(config-router)# network 3.3.3.3/32 area 0.0.0.1	Advertise network 3.3.3.3/32 in OSPF area 1.
R3(config-router)# network 20.1.1.0/24 area 0.0.0.1	Advertise network 20.1.1.0/24 in OSPF area 1.
R3(config-router)#exit	Exit from Router OSPF mode
R3(config)#router bgp 100	Enter Router mode for BGP
R3(config-router)# neighbor 1.1.1.1 remote-as 100	Add neighbor 1.1.1.1 to remote-as 100.
R3(config-router)# neighbor 1.1.1.1 fall-over bfd multihop	Enable BFD option for neighbor
R3(config-router)#neighbor 1.1.1.1 update-source lo	Add neighbor 1.1.1.1 to update-source lo.
R3(config-router)# neighbor 20.1.1.1 remote-as 100	Add neighbor 20.1.1.1 to remote-as 100.
R3(config-router)#neighbor 20.1.1.1 fall-over bfd	Enable BFD option for neighbor
R3(config-router)#commit	Commit the configuration on the node
R3(config-router)#end	Exit from Router BGP mode

Validation

R1

```
R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface   Instance ID
2.2.2.2           1   Full/Backup      00:00:35    10.1.1.2            eth1          0
R1#show bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
=====
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface  Down-
Reason  Remote-Addr
1                   1              IPv4        Single-Hop     Up         00:02:55
eth1      NA                  10.1.1.2/32
```

```
257                 0               IPv4        Multi-Hop      Up          00:00:18
NA          NA                    3.3.3.3/32
Number of Sessions:    2


R1#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
====

Session Interface Index : 10023          Interface name :eth1
Session Index : 1
Lower Layer : IPv4                              Version : 1
Session Type : Single Hop                    Session State : Up
Local Discriminator : 1                         Local Address : 10.1.1.1/32
Remote Discriminator : 1                     Remote Address : 10.1.1.2/32
Local Port : 49152                              Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                    Multiplier: 3
Neg Tx: 250                    Neg Rx: 250                    Neg detect mult: 3
Min echo Tx: 1000          Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000814                Pkt Out : 00000000000000000814
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                  UPTIME : 00:02:59

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 0                   Session Index : 257
Lower Layer : IPv4                               Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 257                     Local Address : 1.1.1.1/32
Remote Discriminator : 0                     Remote Address : 3.3.3.3/32
Local Port : 49153                                Remote Port : 4784
Options :
```

```
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                   Min Rx: 250                   Multiplier: 3
Neg Tx: 0                       Neg Rx: 0                       Neg detect mult: 0
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                Pkt Out : 00000000000000000107
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 0                                              UPTIME : 00:00:00
```

R2

```
R2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface  Instance ID
1.1.1.1           1   Full/DR         00:00:32      10.1.1.1          eth1         0
3.3.3.3           1   Full/Backup     00:00:40     20.1.1.2           eth2         0


R2#show bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
=====
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface   Down-
Reason  Remote-Addr
1               1               IPv4            Single-Hop    Up         00:02:33
eth1            NA              10.1.1.1/32
2               1               IPv4            Single-Hop    Up         00:01:54
eth2            NA              20.1.1.2/32
Number of Sessions:    2


R2#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
====

Session Interface Index : 10023          Interface name :eth1
Session Index : 1
```

```
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                     Local Address : 10.1.1.2/32
Remote Discriminator : 1                 Remote Address : 10.1.1.1/32
Local Port : 49152                          Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                  Min Rx: 250                  Multiplier: 3
Neg Tx: 250                  Neg Rx: 250                  Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000713               Pkt Out : 00000000000000000714
Pkts Drop : 00000000000000000000          Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                UPTIME : 00:02:37

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 10053          Interface name :eth2
Session Index : 2
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                     Local Address : 20.1.1.1/32
Remote Discriminator : 1                 Remote Address : 20.1.1.2/32
Local Port : 49153                           Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                  Min Rx: 250                   Multiplier: 3
Neg Tx: 250                  Neg Rx: 250                  Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled
```

```
Counters values:
Pkt In : 00000000000000000535                Pkt Out : 00000000000000000537
Pkts Drop : 00000000000000000000             Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000              IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                 UPTIME : 00:01:58

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

R3

```
R3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface   Instance ID
2.2.2.2           1     Full/DR         00:00:36    20.1.1.1          eth1            0


R3#show bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
=====
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface   Down-
Reason  Remote-Addr
1                    2          IPv4            Single-Hop      Up         00:01:21
eth1            NA                20.1.1.1/32
257                  1          IPv4            Multi-Hop       Up        00:00:00
NA              NA                1.1.1.1/32
Number of Sessions:    2
Total number of Established sessions 2


R3#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
====

Session Interface Index : 10053          Interface name :eth1
Session Index : 1
Lower Layer : IPv4                            Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                     Local Address : 20.1.1.2/32
Remote Discriminator : 2                 Remote Address : 20.1.1.1/32
Local Port : 49152                            Remote Port : 3784
Options :
```

```
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                    Multiplier: 3
Neg Tx: 250                    Neg Rx: 250                    Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000391                Pkt Out : 00000000000000000391
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                UPTIME : 00:01:26

Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------

Session Interface Index : 0                  Session Index : 257
Lower Layer : IPv4                                Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 257                     Local Address : 3.3.3.3/32
Remote Discriminator : 1                     Remote Address : 1.1.1.1/32
Local Port : 49153                                Remote Port : 4784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                        Min Rx: 250                    Multiplier: 3
Neg Tx: 250                        Neg Rx: 250                    Neg detect mult:250
Min echo Tx: 1000             Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                Pkt Out : 00000000000000000047
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 0                                                UPTIME : 00:00:00
```

```
Protocol Client Info:
BGP-> Client ID: 44      Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

CHAPTER 3 BFD Static Route Configuration

This chapter describes the configurations for BFD static routes.

In order to establish alternate paths to destinations that have the least possible delay it is important to quickly detect any changes to static route validity. BFD detects the liveliness of a static route's nexthop and then uses the nexthop's reachability information to determine whether routes are valid. Using BFD to reach a static route's nexthop also ensures that a static route is inserted in the forwarding database only when the nexthop neighbor is reachable.

Topology

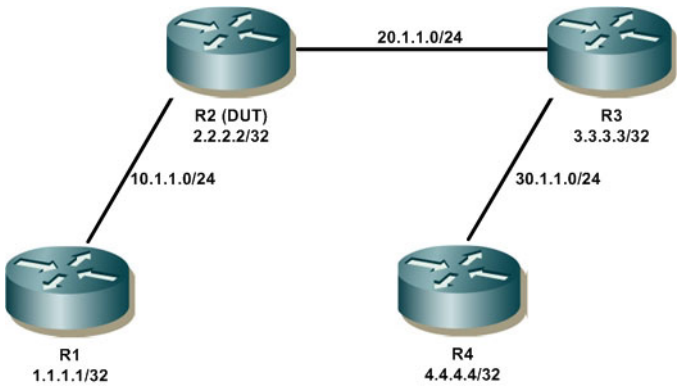

Figure P-11: BFD Static Route Basic Topology

R1

`R1(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address on interface.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit the Interface configuration mode.

R2

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#ip route 30.1.1.0/24 20.1.1.3`	Configure static route.
`R2(config)#ip bfd static all-interfaces`	Enable BFD for all static routes.
`R2(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R2(config-if)#ip static bfd enable`	Enable static BFD on the interface.
`R2(config-if)#ip address 20.1.1.2/24`	Assign IP address on interface.
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exit the Interface configuration mode.
`R2(config)# ip static 30.1.1.0/24 20.1.1.3 fall-over-bfd enable`	Enable static BFD at static route level.
`R2(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.

`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address on interface.
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exit the Interface configuration mode.

R3

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#ip route 10.1.1.0/24 20.1.1.2`	Configure static route.
`R3(config)#ip bfd static all-interfaces`	Enable BFD for all static routes.
`R3(config)#interface eth1`	Enter the Interface configuration mode for `eth1`.
`R3(config-if)#ip address 20.1.1.3/24`	Assign IP address on interface.
`R3(config-if)#ip static bfd enable`	Enable static BFD at interface level.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit the Interface configuration mode.
`R3(config)#ip static 10.1.1.0/24 20.1.1.2 fall-over-bfd enable`	Enable static BFD at static route level.
`R3(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R3(config-if)#ip address 30.1.1.2/24`	Assign IP address on interface.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit the Interface configuration mode.

R4

`R4(config)#interface eth2`	Enter the Interface configuration mode for `eth2`.
`R4(config-if)#ip address 30.1.1.1/24`	Assign IP address on interface.
`R4(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exit the Interface configuration mode.

Validation

R2

```
#sh bfd session
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Remote-
Addr
5          1            IPv4         Single-Hop  Up          00:09:32
20.1.1.3/32
Number of Sessions:    1
VPC2#sh bfd session detail
=======================================================

Session Interface Index : 3              Session Index : 5
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
```

```
Local Discriminator : 5                  Local Address : 20.1.1.2/32
Remote Discriminator : 1                 Remote Address : 20.1.1.3/32
Local Port : 49156                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 2000             Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 0000000000000a29                Pkt Out : 0000000000000bb6
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:09:34

Protocol Client Info:
RIB-> Client ID: 42      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

R3

```
#sh bfd session
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Remote-
Addr
1          5            IPv4         Single-Hop  Up          00:09:39
20.1.1.2/32
Number of Sessions:    1
VPC3#sh bfd session detail
=======================================================

Session Interface Index : 3              Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                  Local Address : 20.1.1.3/32
Remote Discriminator : 5                 Remote Address : 20.1.1.2/32
Local Port : 49152                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess discontinue time : 00:00:00
Bfd GTSM Disabled
Bfd Authentication Disabled
```

```
Counters values:
Pkt In : 0000000000000a59                Pkt Out : 0000000000000a53
Echo Out : 0000000000000000              IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000           IPv6 Pkt Out : 0000000000000000
UP Count : 1                             UPTIME : 00:09:41

Protocol Client Info:
RIB-> Client ID: 42      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

CHAPTER 4 BFD Authentication

This chapter provides BFD authentication configuration examples.

Overview

Bidirectional Forwarding Detection (BFD) is a protocol intended to detect faults in the bidirectional path between two forwarding engines, including physical interfaces, sub-interfaces, and data link. It operates independently of media, data protocols, and routing protocols. A session will be created between links. When BFD links is hacked, the link may be falsely declared to be down, or falsely declared to be up. To overcome this type of situations, need to use authentication on BFD. Using this we can mitigate threats from attackers.

OcNOS supports the following authentication types:

- Simple password
- Meticulous-Keyed-SHA1
- Keyed-SHA1

Among these types, Meticulous-Keyed-SHA1 is the strongest one.

Authentication is optionally enabled on BFD sessions. By default, it is disabled and is configurable via CLI. When authentication is enabled, BFD packets will exchange with authentication section (based on the configured auth type). Receiving system will examine the authentication section of the packet; if it is successful then it will accept. Otherwise, it will discard.

Enabling BFD Authentication for Multihop for IPv4, Using Key-ID

In this example, the BFD Multihop session is configured between R1 and R3 using single key (using key-ID). Once the BFD session is up, the authentication is enabled on both the routers, with the authentication type as Keyed-SHA1. We can enable the authentication on BFD session using any one of the above mentioned authentication type, with the Identical authentication type on both side.

Note: BFD authentication is not supported for Single hop BFD session.

Topology

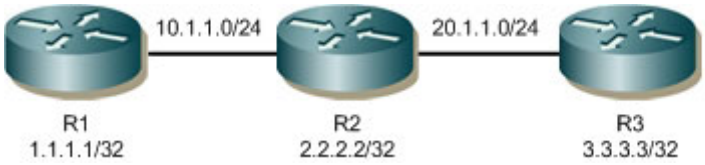

Figure P-12: Basic Topology of Three Routers

Router 1 (R1)

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface lo`	Enter the Loopback Interface configuration mode.
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Assign IP address for interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.

`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#interface xe9`	Enter Interface Mode.
`R1(config-if)#ip address 10.1.1.1/24`	Assign IP address for interface
`R1(config-if)#bfd session 10.1.1.1 20.1.1.2 multihop`	Enabling BFD Multihop session
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit from interface mode.
`R1(config)#router ospf 1`	Enter router OSPF mode
`R1(config-router)#redistribute connected`	
`R1(config-router)#network 10.1.1.0/24 area 0`	Advertise the network to area 0
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exit router OSPF mode
`R1(config)#bfd multihop-peer 20.1.1.2 auth type keyed-sha1 key-id 2 0 key ocnos`	Enabling Authentication for Multihop session

Router 2 (R2)

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface lo`	Enter the Loopback Interface configuration mode.
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Assign IP address for interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#interface xe9`	Enter Interface Mode.
`R2(config-if)#ip address 10.1.1.2/24`	Assign IP address for interface
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#interface xe2`	Enter Interface Mode.
`R2(config-if)#ip address 20.1.1.1/24`	Assign IP address for interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exit from interface mode.
`R2(config)#router ospf 1`	Enter router OSPF mode
`R2(config-router)#redistribute connected`	
`R2(config-router)#network 10.1.1.0/24 area 0`	Advertise the network to area 0
`R2(config-router)#network 20.1.1.0/24 area 0`	Advertise the network to area 0
`R2(config-router)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-router)#exit`	Exit router OSPF mode

Router 3 (R3)

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#interface lo`	Enter the Loopback Interface configuration mode.

R3(config-if)#ip address 3.3.3.3/32 secondary	Assign IP address for interface
R3(config-if)#commit	Commit the candidate configuration to the running configuration.
R3(config-if)#exit	Exit from interface mode.
R3(config)#interface xe2	Enter Interface Mode.
R3(config-if)#ip address 20.1.1.2/24	Assign IP address for interface
R3(config-if)#bfd session 20.1.1.2 10.1.1.1 multihop	Enabling BFD Multihop session
R3(config-if)#commit	Commit the candidate configuration to the running configuration.
R3(config-if)#exit	Exit from interface mode.
R3(config)#router ospf 1	Enter router OSPF mode
R3(config-router)#redistribute connected	
R3(config-router)#network 20.1.1.0/24 area 0	Advertise the network to area 0
R3(config-router)#exit	Exit router OSPF mode
R3(config)#bfd multihop-peer 10.1.1.1 auth type keyed-sha1 key-id 2 0 key ocnos	Enabling Authentication for Multihop session
R3(config)#commit	Commit the candidate configuration to the running configuration.

Validation

Check Multihop session is up between R1 and R3 with authentication type configured.

```
R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/DR         00:00:33    10.1.1.2        xe9
0

R1#sh bfd session

BFD process for VRF: (DEFAULT VRF)
========================================================================
=======
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface       Down-Reason  Remote-Addr
1001       1001         IPv4         Multi-Hop   Up          00:34:32  NA
NA            20.1.1.2/32
Number of Sessions:    1

R1#sh bfd session detail

BFD process for VRF: (DEFAULT VRF)
========================================================================
======

Session Interface Index : 0              Session Index : 1001
Lower Layer : IPv4                       Version : 1
Session Type : Multihop Arbit Path       Session State : Up
Local Discriminator : 1001               Local Address : 10.1.1.1/32
```

```
Remote Discriminator : 1001              Remote Address : 20.1.1.2/32
Local Port : 49152                       Remote Port : 4784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Enabled
Authentication type : keyed-sha1
Authentication Key-id : 2

Counters values:
Pkt In : 00000000000000027907              Pkt Out : 00000000000000028347
Pkts Drop : 00000000000000000000           Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 12                              UPTIME : 00:34:34

Protocol Client Info:
BFD-> Client ID: 28  Flags: 4
------------------------------------------------------------
Number of Sessions:    1

R3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/Backup     00:00:32    20.1.1.1        xe2
0

R3#sh bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
=======
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
1001       1001         IPv4         Multi-Hop   Up          00:36:27  NA
NA           10.1.1.1/32
Number of Sessions:    1

R3#sh bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
======

Session Interface Index : 0              Session Index : 1001
```

```
Lower Layer : IPv4                          Version : 1
Session Type : Multihop Arbit Path          Session State : Up
Local Discriminator : 1001                  Local Address : 20.1.1.2/32
Remote Discriminator : 1001                 Remote Address : 10.1.1.1/32
Local Port : 49152                          Remote Port : 4784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Enabled
Authentication type : keyed-sha1
Authentication Key-id : 2

Counters values:
Pkt In : 00000000000000028428              Pkt Out : 00000000000000028715
Pkts Drop : 0000000000000000000            Auth Pkts Drop :
00000000000000000000
Echo Out : 0000000000000000000             IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 0000000000000000000          IPv6 Pkt Out : 0000000000000000000
UP Count : 12                              UPTIME : 00:36:29

Protocol Client Info:
BFD-> Client ID: 28  Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

CHAPTER 5 BFD with VRF Configuration

This chapter shows using BFD with user defined VRF for OSPFv2 and OSPFv3.

Topology

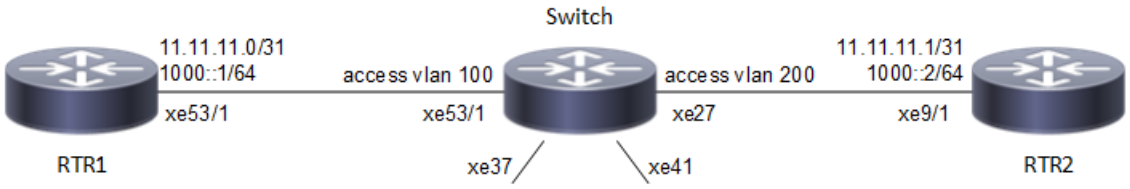

Figure P-13: BFD user-defined VRF

RTR1

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf10`	Create VRF instance with VRF10 and enter into VRF mode
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#interface xe53/1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf10`	Enable VRF forwarding on interface
`(config-if)#ip address 11.11.11.0/31`	Assign IPv4 address in /31 mask.
`(config-if)#ipv6 address 1000::1/64`	Assign IPv6 address in /64 mask.
`(config-if)#ipv6 router ospf area 0 tag vrf10`	Tag OSPFv3 instance on interface with VRF10 for area 0
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 65535 vrf10`	Enter into OSPF VRF configuration mode
`(config-router)#router-id 1.1.1.1`	Assign router ID 1.1.1.1 for OSPF process 65535
`(config-router)#bfd all-interfaces`	Enable BFD for all-interface into OSPF
`(config-router)#network 11.11.11.0/31 area 0`	Enable routing for network 11.11.11.0/31 in area 0
`(config-router)#exit`	Exit from OSPF VRF configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ipv6 vrf ospf vrf10`	Enter into IPv6 OSPF VRF configuration mode
`(config-router)#router-id 2.2.2.2`	Assign router-id 2.2.2.2 for IPv6 OSPF process
`(config-router)#bfd all-interfaces`	Enable BFD for all-interface into OSPF
`(config-router)#exit`	Exit from IPv6 OSPF VRF configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

© 2025 IP Infusion Inc. Proprietary

Switch

`#configure terminal`	Enter configure mode.
`(config)#bridge 32 protocol rstp vlan-bridge`	Create a RSTP VLAN aware bridge with bridge-id 32
`(config)#vlan 100 bridge 32`	Create VLAN 100 and map it to bridge 32
`(config)#vlan 200 bridge 32`	Create VLAN 200 and map it to bridge 32
`(config)#interface xe53/1`	Enter interface mode
`(config-if)#switchport`	Make it port as L2 interface
`(config-if)#bridge-group 32`	Configure bridge group to l2 interface
`(config-if)#switchport mode trunk`	Set the layer 2 interface as trunk interface
`(config-if)#switchport trunk allowed vlan all`	Allow all VLAN to trunk interface
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#interface xe27`	Enter interface mode
`(config-if)#switchport`	Make it port as L2 interface
`(config-if)#bridge-group 32`	Configure bridge group to l2 interface
`(config-if)#switchport mode trunk`	Set the layer 2 interface as trunk interface
`(config-if)#switchport trunk allowed vlan all`	Allow all VLAN to trunk interface
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe37`	Enter interface mode
`(config-if)#switchport`	Make it port as L2 interface
`(config-if)#bridge-group 32`	Configure bridge group to l2 interface
`(config-if)#switchport mode trunk`	Set the layer 2 interface as trunk interface
`(config-if)#switchport trunk allowed vlan all`	Allow all VLAN to trunk interface
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#interface xe41`	Enter interface mode
`(config-if)#switchport`	Make it port as L2 interface
`(config-if)# bridge-group 32 spanning-tree disable`	Configure bridge group to l2 interface with spanning-tree disable
`(config-if)#switchport mode trunk`	Allow all VLAN to trunk interface
`(config-if)#switchport trunk allowed vlan all`	Map the VLAN 200 to access interface
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.

RTR2

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf10`	Create VRF instance with VRF10 and enter into vrf mode
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#interface xe9/1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf10`	Enable VRF forwarding on interface
`(config-if)#ip address 11.11.11.1/31`	Assign IPv4 address in /31 mask.
`(config-if)#ipv6 address 1000::2/64`	Assign IPv6 address in /64 mask.
`(config-if)#ipv6 router ospf area 0 tag vrf10`	Tag OSPFv3 instance on interface with VRF10 for area 0
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 65535 vrf10`	Enter into OSPF VRF configuration mode
`(config-router)#router-id 3.3.3.3`	Assign router ID 3.3.3.3 for OSPF process 65535
`(config-router)#bfd all-interfaces`	Enable BFD for all-interface into OSPF
`(config-router)#network 11.11.11.0/31 area 0`	Enable routing for network 11.11.11.0/31 in area 0
`(config-router)#exit`	Exit from OSPF VRF configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ipv6 vrf ospf vrf10`	Enter into IPv6 OSPF VRF configuration mode
`(config-router)#router-id 4.4.4.4`	Assign router-id 4.4.4.4 for IPv6 OSPF process
`(config-router)#bfd all-interfaces`	Enable BFD for all-interface into OSPF
`(config-router)#exit`	Exit from IPv6 OSPF VRF configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

Validation

RTR1

```
#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 65535 VRF(vrf10):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
3.3.3.3           1   Full/Backup     00:00:32    11.11.11.1      xe15
0

#sh bfd session vrf vrf10

BFD process for VRF: vrf10
========================================================================
=======
```

```
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
256       256         IPv6        Single-Hop  Up         00:26:26  xe15
NA           fe80::eac5:7aff:fe64:4a1d/128
1         1           IPv4        Single-Hop  Up         00:19:05  xe15
NA           11.11.11.1/32
Number of Sessions:    2

#sh bfd session vrf vrf10 detail

BFD process for VRF: vrf10
================================================================================
======

Session Interface Index : 10016          Interface name :xe15
Session Index : 256
Lower Layer : IPv6                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 256                Local Address :
fe80::eac5:7aff:fe78:711d/128
Remote Discriminator : 256               Remote Address :
fe80::eac5:7aff:fe64:4a1d/128
Local Port : 49153                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                               Pkt Out : n/a
Pkts Drop : 00000000000000000000           Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000007256         IPv6 Pkt Out : 00000000000000007256
UP Count : 1                               UPTIME : 00:26:28

Protocol Client Info:
OSPF6-> Client ID: 5  Flags: 4
------------------------------------------------------------

Session Interface Index : 10016          Interface name :xe15
Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                  Local Address : 11.11.11.0/32
Remote Discriminator : 1                 Remote Address : 11.11.11.1/32
Local Port : 49154                       Remote Port : 3784
Options :
```

```
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                                  Pkt Out : n/a
Pkts Drop : 00000000000000000000              Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000               IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000            IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                  UPTIME : 00:19:07

Protocol Client Info:
OSPF-> Client ID: 4  Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

RTR3

```
#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 65535 VRF(vrf10):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/DR         00:00:31    11.11.11.0      xe10
0

#sh ipv6 ospf neighbor

Total number of full neighbors: 1
OSPFv3 Process (vrf10)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR         00:00:28    xe10       0

#sh bfd session vrf vrf10

BFD process for VRF: vrf10
================================================================================
=======
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
1          1            IPv4         Single-Hop  Up          00:20:44  xe10
NA           11.11.11.0/32
256        256          IPv6         Single-Hop  Up          00:28:05  xe10
NA           fe80::eac5:7aff:fe78:711d/128
Number of Sessions:    2

#sh bfd session vrf vrf10 detail
```

```
BFD process for VRF: vrf10
================================================================================
======

Session Interface Index : 10011          Interface name :xe10
Session Index : 1
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                  Local Address : 11.11.11.1/32
Remote Discriminator : 1                 Remote Address : 11.11.11.0/32
Local Port : 49152                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                               Pkt Out : n/a
Pkts Drop : 00000000000000000000           Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 2                               UPTIME : 00:20:46

Protocol Client Info:
OSPF-> Client ID: 4  Flags: 4
------------------------------------------------------------

Session Interface Index : 10011          Interface name :xe10
Session Index : 256
Lower Layer : IPv6                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 256                Local Address :
fe80::eac5:7aff:fe64:4a1d/128
Remote Discriminator : 256               Remote Address :
fe80::eac5:7aff:fe78:711d/128
Local Port : 49153                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
```

```
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                                  Pkt Out : n/a
Pkts Drop : 00000000000000000000              Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000               IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000007707            IPv6 Pkt Out : 00000000000000007718
UP Count : 1                                  UPTIME : 00:28:07

Protocol Client Info:
OSPF6-> Client ID: 5  Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

BFD Over Static Routing IPv4 and IPv6

RTR1

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf10`	Create VRF instance with VRF10 and enter into VRF mode
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#ip bfd vrf vrf10 static all-interfaces`	Enable global IPv4 BFD config for user defined VRF interfaces
`(config)#ipv6 bfd vrf vrf10 static all-interfaces`	Enable global IPv6 BFD config for user defined VRF interfaces
`(config)#interface xe53/1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf10`	Enable VRF forwarding on interface
`(config-if)#ip address 11.11.11.0/31`	Assign IPv4 address in /31 mask.
`(config-if)#ipv6 address 1000::1/64`	Assign IPv6 address in /64 mask.
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#ip route vrf vrf10 200.200.200.200/32 11.11.11.1 xe53/1`	Create IPv4 VRF static route for static BFD session
`(config)#ipv6 route vrf vrf10 2000::/64 1000::2 xe53/1`	Create IPv6 VRF static route for static BFD session
`(config)#exit`	Exit from configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

RTR2

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf10`	Create VRF instance with VRF10 and enter into VRF mode
`(config-vrf)#exit`	Exit from VRF mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#ip bfd vrf vrf10 static all-interfaces`	Enable global IPv4 BFD config for user defined VRF interfaces
`(config)#ipv6 bfd vrf vrf10 static all-interfaces`	Enable global IPv6 BFD config for user defined VRF interfaces
`(config)#interface xe9/1`	Enter interface mode for xe9/1
`(config-if)#ip vrf forwarding vrf10`	Enable VRF forwarding on interface
`(config-if)#ip address 11.11.11.1/31`	Assign IPv4 address in /31 mask.
`(config-if)#ipv6 address 1000::2/64`	Assign IPv6 address in /64 mask.
`(config-if)#exit`	Exit interface mode.

(config)#commit	Commit the candidate configuration to the running configuration.
(config)# ip route vrf vrf10 100.100.100.100/32 11.11.11.0 xe9/1	Create IPv4 VRF static route for static BFD session
(config)# ipv6 route vrf vrf10 3000::2/64 1000::1 xe9/1	Create IPv6 VRF static route for static BFD session
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#exit	Exit from configuration mode

Validation

RTR1

```
#show bfd session vrf vrf10
 BFD process for VRF: vrf10
================================================================================
=============
Sess-Idx  Remote-Disc  Lower-Layer  Sess-Type  Sess-State  UP-Time  Inter-
face  Down-Reason  Remote-Addr
3                   4            IPv4              Single-Hop      Up
00:14:13    xe53/ 1       NA            11.11.11.1/32
4                   5            IPv6              Single-Hop      Up
00:13:24    xe53/ 1       NA             1000::2/128
Number of Sessions:    2


R1#show bfd session vrf vrf10 detail
BFD process for VRF: vrf10
================================================================================
==========

Session Interface Index : 4              Interface name :xe53/1
Session Index : 1
Lower Layer : IPv4                             Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1                      Local Address : 11.11.11.0/32
Remote Discriminator : 1                 Remote Address : 11.11.11.1/32
Local Port : 49152                             Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                    Multiplier: 3
Neg Tx: 250                     Neg Rx: 250                     Neg detect
mult: 3
Min echo Tx: 1000          Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
```

```
Pkt In : 00000000000000001184                    Pkt Out :
00000000000000001184
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000          IPv6 Pkt Out :
00000000000000000000
UP Count : 1                                                    UPTIME :
00:00:21

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------------

Session Interface Index : 4              Interface name :xe53/1
Session Index : 2
Lower Layer : IPv6                             Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                      Local Address :
fe80::5054:ff:fed0:e0c1/128
Remote Discriminator : 2                 Remote Address :
fe80::5054:ff:fe3d:af2b/128
Local Port : 49153                             Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                      Min Rx: 250                    Multiplier: 3
Neg Tx: 250                     Neg Rx: 250                      Neg detect
mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                    Pkt Out :
00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000001158         IPv6 Pkt Out :
00000000000000001163
UP Count : 1                                    UPTIME : 00:00:14

Protocol Client Info:
OSPF6-> Client ID: 5     Flags: 4
------------------------------------------------------------
Number of Sessions:    2


R1#show bfd
BFD ID: 00      Start Time:Tue Mar 19 13:31:08 2019
 BFD Admin State: UP
Number of Sessions:    2
Slow Timer: 2000                      Image type: MONOLITHIC
```

```
Echo Mode: Disabled                BFD Notifications disabled
Next Session Discriminator:    3
```

RTR2

```
#show bfd session vrf vrf10
BFD process for VRF: vrf10
================================================================================
=======
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface  Down-Reason  Remote-Addr
4                  3                          IPv4              Single-Hop
Up       00:15:39    xe9/1       NA            11.11.11.0/32
5                  4                          IPv6              Single-Hop
Up          00:14:50      xe9/1          NA                1000::1/128
Number of Sessions:    2



R2#show bfd session vrf vrf10 detail
BFD process for VRF: vrf10
================================================================================
=======
Session Interface Index : 3               Interface name :xe9/1
Session Index : 1
Lower Layer : IPv4                              Version : 1
Session Type : Single Hop                 Session State : Up
Local Discriminator : 1                      Local Address : 11.11.11.1/32
Remote Discriminator : 1                  Remote Address : 11.11.11.0/32
Local Port : 49152                              Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                     Min Rx: 250                     Multiplier: 3
Neg Tx: 250                     Neg Rx: 250                     Neg detect
mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000001661                   Pkt Out : 00000000000000001665
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                                    UPTIME :
00:00:53

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
------------------------------------------------------------

Session Interface Index : 3               Interface name :xe9/1
```

```
Session Index : 2
Lower Layer : IPv6                              Version : 1
Session Type : Single Hop                  Session State : Up
Local Discriminator : 2                        Local Address :
fe80::5054:ff:fe3d:af2b/128
Remote Discriminator : 2                   Remote Address :
fe80::5054:ff:fed0:e0c1/128
Local Port : 49153                              Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250                    Min Rx: 250                     Multiplier: 3
Neg Tx: 250                     Neg Rx: 250                    Neg detect
mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000000000                  Pkt Out : 00000000000000000000
Echo Out : 00000000000000000000           IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000001642        IPv6 Pkt Out : 00000000000000001636
UP Count : 1                                                  UPTIME :
00:00:56

Protocol Client Info:
OSPF6-> Client ID: 5     Flags: 4
------------------------------------------------------------
Number of Sessions:    2
```

CHAPTER 6 BFD Support on LAG Interface

Overview

The Bidirectional Forwarding Detection (BFD) protocol detects faults in the bidirectional path between two forwarding routers, including interfaces, data links, and beyond. It provides the fastest detection of communication failures on any data links, and the protocol runs over any media and at any protocol layer.

Establish a BFD session using the Link Aggregation Group (LAG) members to identify faults in the bidirectional pathway between two forwarding engines, including interfaces. A single micro-BFD session operates on each member link of the LAG for every enabled address family; that is, each session established for every LAG member link under LAG. The micro-BFD sessions on the member links are individual BFD sessions and operate independently. They use their own distinct local discriminator values, maintain separate sets of state variables, and operate with independent state machines. If a micro-BFD session goes down, the corresponding member link MUST be removed from the LAG load-balancing table.

Single Hop BFD Option Over LAG

To interop with old/legacy routers, where micro-bfd support is not available, and to run regular Single HOP BFD sessions over LAG, a new CLI "`bfd session software`" is being introduced. This software is operated by the control plane, which means the BFD packet TX/RX and the state machine run in the control plane. This session uses a unique discriminator for the whole LAG (not per member interface). BFD Single HOP packets are sent on each member links of LAG on a round-robin basis.

Note:

- When setting up the session with LAG using the "bfd session software" command, the session type will be single-hop.
- When setting up the session with LAG without using the command "bfd session software" command, the session type will Micro- BFD.

Feature Characteristics

BFD is used to control the load-balancing, this protocol minimizes the need for additional MAC addresses, thereby decreasing the resources needed on the Ethernet hardware of the receiving member link. Standard L3 protocols such as OSPF lack visibility into the LAG and perceive it as a single larger interface.

Benefits

BFD offers an efficient way to detect failures in the forwarding path between adjacent routers with low overhead. This encompasses identifying failures in interfaces, data links, and forwarding planes. Enable BFD at the routing protocol level. The BFD protocol identifies faults in the bidirectional path between two forwarding interfaces, encompassing interfaces and data link(s).

Prerequisites

LAG configurations must consist of a minimum of two ports:

- All ports within the LAG must inherit the port characteristics from the primary port.

- For ports within an LAG to ensure a specific port speed, auto-negotiation must be disabled or set to a limited mode.

Additionally, ports in a LAG must be configured for entire duplex operation.

Configuration

The following configuration enables the BFD support on LAG interfaces.

Note: The following configurations are applicable to protocols such as OSPF, BGP, and IS-IS, ensuring immediate notification in the event of session failure.

Topology

The following topology illustrates BFD support on the LAG interface configured between R1, R2, and R3.

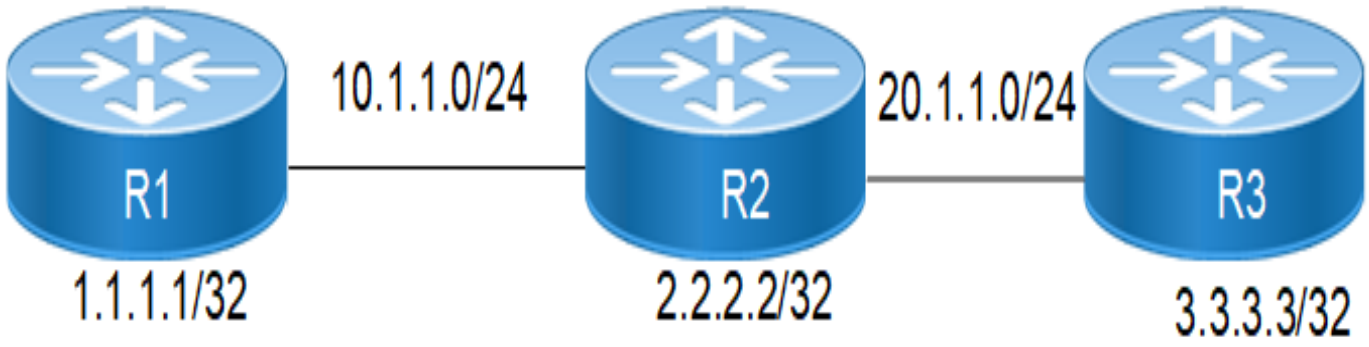

Figure 6-1: BFD support on LAG interface

Perform the following steps to configure the BFD on the LAG interface.

1. Configure the R1 interface with the IP address and with BFD interface.

```
R1#configure terminal
R1(config)#int po1
R1(config-if)#ip address 10.1.1.2/24
R1(config)#commit
R1(config)#interface eth1
R1(config-if)#channel-group 1 mode active
R1(config)#commit
R1#configure terminal
R1(config-if)#bfd session 10.1.1.2 10.1.1.1
R1(config-if)#bfd session ?
A.B.C.D   Source IPv4 address
X:X::X:X  Source IPv6 address
hardware  BFD Session offloaded to Forwarding Plane
software  BFD Session running at Control Plane (Host CPU)
R1(config-if)#bfd session software
R1(config)#commit
```

2. Configure the R2 interface with the IP address and with BFD interface.

```
R2#configure terminal
```

```
R2(config)#int po1
R2(config-if)#ip address 10.1.1.1/24
R2(config)#commit
R2(config)#interface eth1
R2(config-if)#channel-group 1 mode active
R2(config)#commit
R2#configure terminal
R2(config-if)#bfd session 10.1.1.1 10.1.1.2
R2(config-if)#bfd session ?
A.B.C.D   Source IPv4 address
X:X::X:X  Source IPv6 address
hardware  BFD Session offloaded to Forwarding Plane
software  BFD Session running at Control Plane (Host CPU)
R2(config-if)#bfd session software
R2(config)#commit
```

Validation

R1

The following show output displays the BFD support on LAG interface.

```
OcNOS#show bfd session

BFD process for VRF: (DEFAULT VRF)
========================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
257       1002        IPv4        Single-Hop  Up         00:05:49  po1
NA           10.1.1.2/32

Number of Sessions:    1
OcNOS#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
========================================================================

Session Interface Index : 100001         Interface name :po1
Session Index : 257
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1002               Local Address : 10.1.1.1/32
Remote Discriminator : 1002              Remote Address : 10.1.1.2/32
Local Port : 49156                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250              Multiplier: 3
Neg Tx: 250               Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
```

```
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : 00000000000000001255                Pkt Out : 00000000000000001291
Echo Out : 00000000000000000000              IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000           IPv6 Pkt Out : 00000000000000000000
UP Count : 2                                 UPTIME : 00:04:39

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------------

Number of Sessions:    1
```

R2

The following show output displays the BFD support on LAG interface.

```
 OcNOS#show bfd session

BFD process for VRF: (DEFAULT VRF)
================================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
8196      1002        IPv4        Single-Hop  Up         00:09:50  po1
NA           10.1.1.1/32

Number of Sessions:    1
OcNOS#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
================================================================================
Session Interface Index : 100001         Interface name :po1
Session Index : 8196
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 1002               Local Address : 10.1.1.2/32
Remote Discriminator : 1002              Remote Address : 10.1.1.1/32
Local Port : 49155                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                             Pkt Out : n/a
```

```
Echo Out : 00000000000000000000            IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 2                               UPTIME : 00:10:11

Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
-------------------------------------------------------------

Number of Sessions:    1
```

New CLI Commands

bfd session software

Use this command to create a BFD session operated by control plane (This implies that the BFD TX or RX will be operated in the control plane).

Note: The scaling number of BFD sessions depends on the CPU performance of the target device since a Single Hop BFD session over LAG is handled in the control plane. Therefore, it is recommended to limit the number of sessions to a maximum of 8, with a minimum BFD interval of 100ms (Default BFD interval is 250ms). The lowest configurable interval for a software-based BFD session is 50ms.

Command Syntax

```
bfd session software
```

Parameters

Not applicable

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 6.5.1.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#bfd session software
(config-if)#
```

Glossary

Key Terms/Acronym	Description
BFD	Bidirectional Forwarding Detection
LAG	Link Aggregation Group

OSPF	Open Shortest Path First
MAC	Media Access Control Address.

Bidirectional Forwarding Detection Command Reference

© 2025 IP Infusion Inc. Proprietary

CHAPTER 1 Bidirectional Forwarding Commands

This chapter explains the commands used to configure Bidirectional Forwarding (BFD):

- accept-lifetime
- bfd auth type
- bfd
- bfd echo
- bfd echo interval
- bfd echo ipv4 source
- bfd interval
- bfd multihop-peer
- bfd multihop-peer A.B.C.D interval
- bfd multihop-peer interval
- bfd multihop-peer X:X::X:X interval
- bfd notification
- bfd session
- bfd slow-timer
- debug bfd
- key
- key chain
- key-string
- key-string encrypted
- send-lifetime
- show bfd
- show bfd interface
- show bfd session
- show bfd session A.B.C.D
- show bfd session ipv6
- show debugging bfd
- snmp restart bfd

© 2025 IP Infusion Inc. Proprietary

accept-lifetime

Use this command to specify the time period during which the authentication on a key chain is received as valid.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> HH:MM:SS MONTH <01-31> <1993-2035>
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> infinite
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> duration <1-2147483646>
no accept-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of accept-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to start. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to start as the first three letters of the month with first letter in upper case, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the end time of accept-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to end. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to end as the first three letters of the month with first letter in caps, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to end.
`duration`	Indicate the duration parameter.
`<1-2147483646>`	Specify the actual end time duration of a key in seconds.
`infinite`	Specify the end time to never expire.

Default

By default, accept-lifetime command is disabled

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of accept-lifetime for key1 on the key chain named mychain.

```
#configure terminal
(config)#key chain mychain
```

```
(config-keychain)#key 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 03 2004 04:04:02 Oct 06
2006
```

bfd auth type

Use this command to enable BFD authentication on an interface that has either an IPv4 or an IPv6 BFD session.

Use the `no` form of the command to disable BFD authentication.

Note: BFD authentication is not supported for single hop IPv4 sessions.

Command Syntax

```
bfd auth type (keyed-md5 | keyed-sha1 | meticulous-keyed-md5 | meticulous-keyed-
  sha1 | simple) key-id <0-255> <0|1> key (WORD)
no bfd auth
```

Parameters

`auth type`	Specify an authentication type.
`keyed-md5`	Specify a keyed message digest authentication type.
`simple`	Specify a simple authentication type.
`keyed-sha1`	Specify a keyed secure hashing algorithm authentication type.
`meticulous-keyed-md5`	Specify an authentication key meticulous keyed message digest authentication.
`meticulous-keyed-sha1`	Specify an authentication key meticulous keyed secure hashing algorithm authentication.
`key-id`	Indicate the `key-id` keyword.
`<0-255>`	Specify the key ID value.
`<0|1>`	0 Unencrypted password (key) 1 Encrypted password (key)
`key`	Indicate the `key` keyword.
`WORD`	Specify the authentication key name. Whitespaces is not allowed in key.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Do the following to configure a single-key support:

Use value 0 before the key parameter for plain text key.

Use value 1 before the key parameter for already encrypted key.

```
(config)#interface eth1
(config-if)#bfd auth type keyed-md5 key-id 10 0 key test-key
(config-if)#bfd auth type keyed-sha1 key-id 20 1 key 0x4d8ea8d7e9ee090b
```

bfd

Use this command to enable and disable all the BFD sessions on this interface.

Command Syntax

```
bfd (enable|disable)
```

Parameters

`enable`	Enable BFD
`disable`	Disable BFD

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#bfd disable
```

bfd echo

Use this command to set BFD sessions to echo mode.

Use the `no` form of the command to return a BFD session to its default mode.

Note: BFD echo mode is supported for IPv4 BFD single hop sessions only.

Command Syntax

```
bfd echo
no bfd echo
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd echo
```

bfd echo interval

Use this command to set the BFD echo interval.

Use the `no` form of this command to reset the echo interval to its default value.

Command Syntax

```
bfd echo interval <50-4294967>
no bfd echo interval
```

Parameter

`<50-4294967>`	Transmit interval in milliseconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following command sets the BFD echo with no values.

```
#configure terminal
(config)#interface eth1
(config-if)#bfd echo interval 234
```

bfd echo ipv4 source

Use this command to set the echo ipv4 source address.

Use the `no` form of this command to reset the value.

Command Syntax

```
bfd echo ipv4 source A.B.C.D
no bfd echo ipv4 source
```

Parameter

`A.B.C.D`	Source IPv4 address.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following command displays the ipv4 source address value.

```
#configure terminal
(config)#interface eth1
(config-if)#bfd echo interval 234
(config-if)#bfd echo ipv4 source A.B.C.D
```

bfd interval

Use this command to configure BFD transmit and receive intervals, and the hello multiplier value.

Use the `no` form of the command to set the intervals and multiplier to their default values.

Note: If “bfd session software” configured on any l3 interface then minimum tx/rx can be configured as 10 else minimum tx/rx can be configured as 3.

Command Syntax

```
bfd interval <3-999> minrx <3-999> multiplier <3-50>
no bfd interval
```

Parameters

`<3-999>`	Transmit interval in milliseconds.
`minrx`	Receive interval.
`<3-999>`	Receive interval in milliseconds.
`multiplier`	Hello multiplier.
`<3-50>`	Hello multiplier value.

Defaults

The default for the transmit and receive intervals is 250 milliseconds.

The default hello multiplier value is 3.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#bfd interval 100 minrx 100 multiplier 5
(config-if)#
```

bfd multihop-peer

Use this command to enable authentication over either a multihop IPv4 or IPv6 session

Use the `no` form of the command to disable BFD authentication.

Command Syntax

```
bfd multihop-peer (A.B.C.D | X:X::X:X) auth type (keyed-md5 | keyed-sha1 |
  meticulous-keyed-md5 | meticulous-keyed-sha1 | simple) key-id <0-255> <0|1> key
  (WORD)
no bfd multihop-peer (A.B.C.D) auth
```

Parameters

`A.B.C.D`	Specify an IPv4 address.
`X:X::X:X`	Specify
`auth type`	Specify an authentication type.
`simple`	Specify a simple authentication type.
`keyed-md5`	Specify a keyed message digest authentication type.
`keyed-sha1`	Specify a keyed secure hashing algorithm authentication type.
`meticulous-keyed-sha1`	Specify an authentication key Meticulous Keyed Secure hashing algorithm authentication type.
`meticulous-keyed-md5`	Specify an authentication key meticulous keyed message digest authentication.
`key-id`	Indicate the `key-id` keyword.
`<0-255>`	Specify the key ID value.
`<0|1>`	0 Unencrypted password (key) 1 Encrypted password (key)
`key`	Indicate the `key` keyword.
`WORD`	Specify the authentication key name. Whitespaces is not allowed in key.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Do the following to configure a single-key support:

Use value 0 before the key parameter for plain text key.

Use value 1 before the key parameter for already encrypted key.

```
#configure terminal
```

```
(config)#bfd multihop-peer 10.10.10.2 auth type keyed-md5 key-id 100 0 key
test-key
(config)#bfd multihop-peer 20.20.20.2 auth type keyed-sha1 key-id 200 1 key
0x4d8ea8d7e9ee090b
```

bfd multihop-peer A.B.C.D interval

Use this command to configure IPv4 BFD multihop peer timer values.

Use the `no` form of the command to reset the IPv4 multihop peer timer value.

Note: Multihop BFD operates at control plane. When handling numerous services or traffic sessions, BFD session instability might occur. To address this, it's recommended to increase RX/TX interval to the maximum value 999 milliseconds and refrain from selecting or configuring the default 250-millisecond option.

Command Syntax

```
bfd multihop-peer A.B.C.D interval <50-999> minrx <50-999> multiplier <3-50>
no bfd multihop-peer A.B.C.D interval
```

Parameters

`interval`	Indicate the interval parameter.
`<50-999>`	Specify the actual transmit interval in milliseconds.
`minrx`	Indicate the minrx parameter.
`<50-999>`	Specify the actual reception interval in milliseconds.
`multiplier`	Indicate the multiplier parameter.
`<3-50>`	Specify the actual hello multiplier value.

Command Mode

Configure mode

Default

Multiplier value is 3

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd multihop-peer 10.1.1.67 interval 100 minrx 100 multiplier 3
```

Note:

Command Syntax

```
bfd multihop-peer A.B.C.D interval <50-999> minrx <50-999> multiplier <3-50>
no bfd multihop-peer A.B.C.D interval
```

Parameters

`interval`	Indicate the interval parameter.
`<50-999>`	Specify the actual transmit interval in milliseconds.
`minrx`	Indicate the minrx parameter.
`<50-999>`	Specify the actual reception interval in milliseconds.

`multiplier`	Indicate the multiplier parameter.
`<3-50>`	Specify the actual hello multiplier value.

Command Mode

Configure mode

Default

Multiplier value is 3

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd multihop-peer 10.1.1.67 interval 100 minrx 100 multiplier 3
(config)#
```

bfd multihop-peer interval

Use this command for the global configuration of the timers for all multi-hop BFD sessions.

Use the `no` form of the command to reset the multihop peer timer value.

Note: Multihop BFD operates at the control plane. When handling numerous services or traffic sessions, BFD session instability might occur. To address this, the recommendation is to increase the RX/TX interval to the maximum value of 999 milliseconds and refrain from selecting or configuring the default 250-millisecond option.

Command Syntax

```
bfd multihop-peer interval <50-999> minrx <50-999> multiplier <3-50>
no bfd multihop-peer interval
```

Parameters

`interval <50-999>`	Specifies the actual transmit interval in this range of milliseconds. The default transmit interval is 250 milliseconds.
`minrx <50-999>`	Specify the actual reception interval in milliseconds. The default reception interval is 250 milliseconds.
`multiplier <3-50>`	Specify the actual hello multiplier value. The default multiplier value is 3.

Command Mode

Configure mode

Default

Disabled

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example shows how to configure timers for all multi-hop BFD sessions:

```
#configure terminal
(config)#bfd multihop-peer interval 100 minrx 100 multiplier 3
```

bfd multihop-peer X:X::X:X interval

Use this command to configure an IPv6 BFD multihop peer timer values.

Use the `no` form of the command to reset the IPv6 multihop peer timer values.

Command Syntax

```
bfd multihop-peer X:X::X:X interval <50-999> minrx <50-999> multiplier <3-50>
no bfd multihop-peer X:X::X:X interval
```

Parameters

`interval`	Indicate the interval parameter.
`<50-999>`	Specify the actual transmit interval in milliseconds.
`minrx`	Indicate the minrx parameter.
`<50-999>`	Specify the actual reception interval in milliseconds.
`multiplier`	Indicate the multiplier parameter.
`<3-50>`	Specify the actual hello multiplier value.

Command Mode

Configure mode

Default

Default multiplier value is 3

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd multihop-peer 10.1.1.1 interval 100 minrx 100 multiplier 3
```

bfd notification

Use this command to enable or disable BFD notification.

Command Syntax

```
bfd notification (enable | disable)
```

Parameters

`disable`	Disable BFD notification.
`enable`	Enable BFD notification.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#bfd notification enable

(config)#bfd notification disable
```

bfd session

Use this command to create a BFD IPv4 or IPv6 session on an interface.

Use the no form of the command to delete a BFD session from an interface.

Command Syntax

```
interface <IFNAME>/bfd session (A.B.C.D | X:X::X:X) (A.B.C.D | X:X::X:X) (multihop)
  (non-persistent | admin-down | demand-mode)
```

Parameters

`A.B.C.D`	Source IPv4 address
`X:X::X:X`	Source IPv6 address
`multihop`	Multihop Session
`admin-down`	Administrative down session
`demand-mode`	Demand mode session
`non-persistent`	Non persistent session

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
OcNOS(config)#interface eth2
OcNOS(config-if)#bfd session 20.20.20.2 20.20.20.1 admin-down
OcNOS(config-if)#do sh run bfd
!
interface eth2
bfd session  20.20.20.2  20.20.20.1 admin-down
!
OcNOS(config-if)#no bfd session  20.20.20.2  20.20.20.1 admin-down
OcNOS(config-if)#do sh run bfd
!
interface eth2
bfd session  20.20.20.2  20.20.20.1
!
OcNOS(config-if)#bfd session 20.20.20.2 20.20.20.1 admin-down
OcNOS(config-if)#do sh run bfd
!
interface eth2
bfd session  20.20.20.2  20.20.20.1 admin-down
!
OcNOS(config-if)#no bfd session 20.20.20.2 20.20.20.1
OcNOS(config-if)#do sh run bfd
!
```

bfd slow-timer

Use this command to set a BFD slow timer interval.

Use the `no` form of the command to reset the timer to default values.

Command Syntax

```
bfd slow-timer <1000-30000>
no bfd slow-timer
```

Parameter

`<1000-30000>`	Interval for the slow-timer in milliseconds

Command Mode

Configure mode

Default

Default slow-timer value is 2000

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd slow-timer 1500
(config)#
```

debug bfd

Use this command to enable debugging for BFD.

Use the `no` form of the command to disable all debugging for BFD.

Command Syntax

```
debug bfd (all|)
debug bfd (event|ipc-error|ipc-event|nsm|packet|session)
no debug bfd (all|)
no debug all
no debug bfd (event|ipc-error|ipc-event|nsm|packet|session)
```

Parameters

`all`	Enable all debugging.
`event`	Enable BFD event debugging.
`ipc-error`	Enable BFD IPC-error debugging
`ipc-event`	Enable BFD IPC-event debugging.
`nsm`	Enable BFD NSM debugging.
`packet`	Enable BFD packet debugging.
`session`	Enable BFD session debugging.

Command Mode

Exec, Privileged Exec and Configure Mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug bfd all
#debug bfd event
#debug bfd ipc-error
#debug bfd ipc-event
#debug bfd nsm
#debug bfd packet
#debug bfd session
```

key

Use this command to manage, add or delete authentication keys in a key chain. This command allows you to enter the keychain-key mode to set a password for the key.

Command Syntax

```
key <0-2147483647>
no key <0-2147483647>
```

Parameters

`<0-2147483647>` Specify a key identifier.

Default

By default, BFD uses level-1-2 if there is no Level-2 instance nor a Level-1-2 instance. Otherwise, it uses level-1.

Command Mode

Keychain mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example configures a key number 1 and shows the change to keychain-key command mode.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#
```

key chain

Use this command to enter key chain mode to configure a key chain with a key chain name. This command allows you to enter the keychain mode to specify keys on this key chain.

Command Syntax

```
key chain WORD
no key chain WORD
```

Parameters

WORD	Specify the name of the key chain to manage.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the creation of a key chain named `mychain` and the change to keychain mode:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#
```

key-string

Use this command to define a password in plain-text to be used by a key.

The password is stored as encrypted, and is displayed in encrypted text when show running-config command is executed.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
key-string WORD
no key-string
```

Parameters

`WORD`	Specify a string of characters to be used as a password by the key. The length of the string should be between 1-80 characters.

Default

By default, password is not configured.

Command Mode

Keychain mode and Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#key-string prime

(config-keychain)#key 1
(config-keychain-key)#no key-string
```

key-string encrypted

Use this command to define a password in its encrypted format to be used by a key.

Use the `no` parameter with this command to disable this feature

Command Syntax

```
key-string encrypted WORD
no key-string encrypted
```

Parameters

`WORD`	Specify the encrypted string of characters to be used as a passoword by the key. The length of this string should be between 18-162 characters.

Default

By default, password is not configured.

Command Mode

Keychain mode and Keychain-key mode.

Applicability

This command was introduced in OcNOS version 4.1.

Examples:

In the following example, the encrypted password for key 1 in the key chain named mykeychain is set to 0xd6c50b442de47f70 (equivalent to "mychain" in plain-text):

```
#configure terminal
(config)#key chain mykeychain
(config-keychain)#key 1
(config-keychain-key)#key-string encrypted 0xd6c50b442de47f70
(config-keychain)#key 1
(config-keychain-key)#no key-string
```

send-lifetime

Use this command to specify the time period during which the authentication key on a key chain can be sent.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> HH:MM:SS MONTH <01-31> <1993-2035>
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> infinite
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> duration <1-2147483646>
no send-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of send-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to start. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to start as the first three letters of the month with first letter in upper case, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the end time of send-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to end.
`MONTH`	Specify the month of the year to end as the first three letters of the month with first letter in upper case, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to end.
`duration`	Indicate the duration parameter.
`<1-2147483646>`	Specify the actual end time duration of a key in seconds.
`infinite`	Specify the end time to never expire.

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of `send-lifetime` for `key 1` on the key chain named `mychain`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#send-lifetime 03:03:01 Jan 03 2004 04:04:02 Dec 06 2006
```

show bfd

Use this command to display information about the BFD process.

Command Syntax

```
show bfd
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below displays the command syntax and sample output from the command.

```
#show bfd
BFD ID: 00      Start Time:Fri May 1 09:55:06 2009
Number of Sessions:     1
Slow Timer: 1000        Image type: MONOLITHIC
Echo Mode: Disabled     Next Session Discriminator:    2
#
```

Table P-5 explains the output fields.

Table P-5: show BFD fields

Entry	Description
BFD ID	ID number of the BFD session.
Start Time	The date and time when the BFD session was started
BFD Admin State	State of the BFD session: Initializing: session is initializing Up: session is up. Down: session is down
Number of Sessions	Number of BFD sessions running on the device.
Slow Timer	Required minimum transmission time for the BFD session.
Image Type	Distributed or Monolithic.
Echo Mode	Either enabled or disabled.
Next Session Discriminator	An opaque discriminator value that identifies each session on the device that is used to demultiplex multiple BFD sessions between the same pair of devices.

show bfd interface

Use this command to display details for an interface running BFD or for all interfaces configured for BFD.

Command Syntax

```
show bfd interface (ifindex <0-4294967295>|all|)
```

Parameters

`all`	Display all interfaces.
`ifindex`	Display an interface index.
`<0-4294967295>`	Display an ID of an interface in this range.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below displays the command syntax and sample output from the command.

```
#show bfd interface all
Interface:         lo ifindex: 1 state:   UP
Interface level configuration: NO ECHO, NO SLOW TMR
Timers in Milliseconds
Min Tx: 20 Min Rx: 20 Multiplier: 5

Interface:       eth0 ifindex: 2 state:   UP
Interface level configuration: NO ECHO, NO SLOW TMR
Timers in Milliseconds
Min Tx: 20 Min Rx: 20 Multiplier: 5

Interface:       eth1 ifindex: 3 state: DOWN
Interface level configuration: NO ECHO, NO SLOW TMR
Timers in Milliseconds
Min Tx: 20 Min Rx: 20 Multiplier: 5

Interface:       sit0 ifindex: 4 state: DOWN
Interface level configuration: NO ECHO, NO SLOW TMR
Timers in Milliseconds
Min Tx: 20 Min Rx: 20 Multiplier: 5

Interface:       gre0 ifindex: 5 state: DOWN
Interface level configuration: NO ECHO, NO SLOW TMR
Timers in Milliseconds
Min Tx: 20 Min Rx: 20 Multiplier: 5
```

Table P-6 explains the output fields.

Table P-6: show BFD interface fields

Entry	Description
interface	Interface on which BFD is running
ifindex	Interface index number
state	State of the BFD session for the interface: UP: session is up. DOWN: session is down
Interface level configuration	State of interface-level properties: echo on or off slow-timer (timer dampening) on or off
Min Tx	The minimum interval that the local device would like to use when transmitting BFD control packets.
Min Rx	The minimum interval between received BFD control packets that this device is capable of supporting
Multiplier	The negotiated transmit interval, multiplied by this value, provides the detection time for the receiving device in asynchronous mode.

show bfd session

Use this command to display BFD sessions.

Note: BFD Packet In and BFD Packet Out counters are not supported for Single Hop IPv4 BFD sessions on Qumran devices.

Command Syntax

```
show bfd session (detail|)
show bfd session vrf (WORD|all|default) (detail|)
```

Parameters

`detail`	Session details.
`WORD`	Sessions for this user-defined Virtual Routing and Forwarding instance name.
`all`	Show information for all Virtual Routing and Forwarding instances
`default`	Show information for only the default Virtual Routing and Forwarding instance

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show bfd session detail
Session Interface Index : 3                 Session Index : 1
Lower Layer : IPv4                          Version : 1
Session Type : Single Hop                   Session State : Down
Local Discriminator : 1                     Local Address : 19.19.19.2/32
Remote Discriminator : 0                    Remote Address : 19.19.19.1/32
Local Port : 49152                          Remote Port : 3784
Options :

Diagnostics: None

Timers in Milliseconds :
Min Tx: 20              Min Rx: 20              Multiplier: 5
Neg Tx: 0               Neg Rx: 0               Neg detect mult: 0
Min echo Tx: 20         Min echo Rx: 10         Neg echo intrvl: 0
Storage type: 2
Sess down time: 00:00:00
Bfd Authentication Enabled
Authentication type : simple
Auth-Key-Id: 47

Counters values:
Pkt In : 0000000000000000                   Pkt Out : 0000000000000011
Echo Out : 0000000000000000                 IPv6 Echo Out : 0000000000000000
IPv6 Pkt In : 0000000000000000              IPv6 Pkt Out : 0000000000000000
UP Count : 0                                UPTIME : 00:00:00
```

```
Protocol Client Info:
BFD-> Client ID: 28      Flags: 4
------------------------------------------------------


#show bfd session vrf all

BFD process for VRF: vrf1
========================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
1          1           IPv4         Single-Hop  Up          00:05:38  eth1
NA           1.1.1.2/32
Number of Sessions:    1

BFD process for VRF: (DEFAULT VRF)
========================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
1          7           IPv4         Single-Hop  Up          00:03:31  eth3
NA           3.3.3.2/32
2          9           IPv4         Single-Hop  Up          00:00:45  eth4
NA           4.4.4.2/32
Number of Sessions:    2

BFD process for VRF: vrf2
========================================================================
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time
Interface      Down-Reason  Remote-Addr
1          6           IPv4         Single-Hop  Up          00:03:37  eth2
NA           2.2.2.2/32
Number of Sessions:    1
```

Table P-7 explains the output fields.

Table P-7: show bfd session detail fields

Entry	Description
Session Interface Index	ID number of the interface.
Session Index Sess-Idx	ID number of this BFD session.
Lower Layer	The lower layer protocol on which BFD is carried: IPv4 IPv6
Version	Session version number; generally 1.
Session Type Sess-Type	Single Hop Multihop Arbit Path Multihop OOB Signalled Multihop Unidirectional.

Table P-7: show bfd session detail fields (Continued)

Entry	Description
Session State Sess-State	The State of the session: Init: The session is initializing Up: The session is up. Down: The session is down AdminDown: The session has been administratively shutdown indefinitely.
Local Discriminator	A unique discriminator value generated by the transmitting device used to demultiplex multiple BFD sessions between the same pair of devices.
Local Address	Local address of the transmitting device.
Remote Discriminator Remote-Disc	The discriminator received from the corresponding remote device; zero if that value is unknown.
Remote Address Remote-Addr	Remote address of the receiving device.
Local Port	UDP port number of the transmitting device.
Remote Port	UDP port number of destination.
Options	Fate Shared Echo Enabled Demand Enabled Remote Demand Enbl Remote admin Down Poll seq Init.
Diagnostics	Performance diagnostics: None Control Detection Time Expired Echo Failed Neighbor Session Down Forwarding Plane Reset Path Down Concatenated Path Down Admin Down Reverse Concatenated Path Down
Min Tx	Minimum transmit interval.
Min Rx	Minimum receive interval.
Multiplier	The negotiated transmit interval, multiplied by this value, provides the detection time for the receiving system in asynchronous mode.
Neg Tx	Negotiated transmit interval.
Neg Rx	Negotiated receive interval in milliseconds.
Neg detect mult	Negotiated detection multiplier.
Min echo Tx	Minimum echo transmit interval in milliseconds.

Table P-7: show bfd session detail fields (Continued)

Entry	Description
Min echo Rx	Minimum echo receive interval in milliseconds.
Neg echo intrvl	Negotiated echo interval.
Storage type	SNMPv2 storage type (usually be set to 2): other(1) volatile(2) non-Volatile(3) permanent(4 read-Only(5)
Sess down time	Length of time this BFD session has been down.
Bfd GTSM	Whether the BFD session has enabled or disabled the Generalized TTL Security Mechanism (GTSM), which uses the time to live (TTL) or hop count to prevent off-link attackers from spoofing packets.
Bfd Authentication Enabled	When authentication is enabled.
Authentication type	Type of authentication: simple: Simple Password keyed-md5: Keyed message digest keyed-sha1: Keyed Secure hashing algorithm meticulous-keyed-md5: Meticulous Keyed message digest meticulous-keyed-sha1: Meticulous Keyed Secure hashing algorithm
Authentication Key-id	ID number of the authentication key.
Pkt In	Number of packets that have been received during this BFD session.
Pkt Out	Number of packets that have been transmitted during this BFD session.
Echo Out	Number of Echo-packets that have been transmitted during this BFD session.
IPv6 Pkt In	Number of IPv6 packets that have been received during this BFD session.
IPv6 Pkt Out	Number of IPv6 packets that have been transmitted during this BFD session.
UP Count	Number of times the BFD session has been up.
UPTIME	Length of time this BFD session has been up.
BGP-> Client ID	Protocol and module ID number of this BFD session’s neighboring device.
Flags	Session state of the neighboring device.
Interface	The interface on which the VRF resides.
Down-Reason	The reason causing the VRF to be down.

show bfd session A.B.C.D

Use this command to display information about an IPv4 BFD session neighbor.

Command Syntax

```
show bfd session A.B.C.D A.B.C.D (detail|)
show bfd session vrf WORD A.B.C.D A.B.C.D (detail|)
show bfd session A.B.C.D A.B.C.D <0-4294967295> (detail|)
show bfd session vrf WORD A.B.C.D A.B.C.D <0-4294967295> (detail|)
```

Parameters

WORD	Name of a specific Virtual Routing and Forwarding instance
A.B.C.D	Display the local IPv4 address.
A.B.C.D	Display the neighbor IPv4 address.
<0-4294967295>	Display the interface index of the address.
detail	Display detailed information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below displays the command syntax and sample output from the command.

```
#show bfd session 10.1.1.66 10.1.1.67 3
Session Interface Index: 3     Session Index: 1
Lower Layer: IPv4     Single Hop
Session State: Up
Local Discriminator: 1 Remote Discriminator: 163
Local Address: 10.1.1.66/32    Remote Address: 10.1.1.67/32
Local Port: 49152      Remote Port: 3784
Timers in Milliseconds
Min Tx: 1000 Min Rx: 1000 Multiplier: 4
UP Count: 1 UPTIME: 00:10:08


#show bfd session vrf raj 130.1.1.1 130.1.1.2 detail

Session Interface Index : 10017          Interface name :xe16
Session Index : 2
Lower Layer : IPv4                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 2                  Local Address : 130.1.1.1/32
Remote Discriminator : 2                 Remote Address : 130.1.1.2/32
Local Port : 49153                       Remote Port : 3784
Options :
```

```
Diagnostics : None

Timers in Milliseconds :
Min Tx: 250               Min Rx: 250               Multiplier: 3
Neg Tx: 250               Neg Rx: 250               Neg detect mult: 3
Min echo Tx: 1000         Min echo Rx: 1000         Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                                  Pkt Out : n/a
Pkts Drop : 00000000000000000000              Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000               IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000000000            IPv6 Pkt Out : 00000000000000000000
UP Count : 1                                  UPTIME : 00:01:38

Protocol Client Info:
OSPF-> Client ID: 4      Flags: 4
```

Table P-8 explains the output fields.

Table P-8: Show BFD session fields

Entry	Description
Session Interface Index	ID number of the Interface.
Session Index	ID number for this BFD session.
Lower Layer	The lower layer protocol on which BFD is carried (IPv4 or IPv6).
Session State	The state of the session: Init: session is initializing Up: session is up. Down: session is down AdminDown: session has been administratively shutdown indefinitely
Local Discriminator	A unique value generated by the transmitting device used to demultiplex multiple BFD sessions between the same pair of devices.
Remote Discriminator	The discriminator received from the corresponding remote device. This field is zero if the value is unknown.
Local Address	Local address of the transmitting device.
Remote Address	Remote address of the receiving device.
Local Port	UDP port number of the transmitting device.
Remote Port	UDP port number of the receiving device.
Min Tx	Minimum transmit interval in milliseconds.

Table P-8: Show BFD session fields

Entry	Description
Min Rx	Minimum receive interval in milliseconds.
Multiplier	The negotiated transmit interval, multiplied by this value, provides the detection time for the receiving system in asynchronous mode.
UP Count	The number of times the BFD session has been in up state.
UPTIME	The length of time this BFD session has been in the up state.

show bfd session ipv6

Use this command to display information about an IPv6 BFD session neighbor.

Command Syntax

```
show bfd session ipv6 X:X::X:X X:X::X:X (detail|)
show bfd session ipv6 vrf WORD X:X::X:X X:X::X:X (detail|)
show bfd session ipv6 X:X::X:X X:X::X:X <0-4294967295> (detail|)
show bfd session ipv6 vrf WORD X:X::X:X X:X::X:X <0-4294967295> (detail|)
```

Parameters

WORD	Name of a specific Virtual Routing and Forwarding instance
X:X::X:X	Display the local IPv6 address.
X:X::X:X	Display the neighbor IPv6 address.
<0-4294967295>	Display the interface index of the address.
detail	Display detailed information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below displays the command syntax and sample output from the command.

```
#show bfd session 2001::1222 2001::1223 3
Session Interface Index : 3     Session Index: 1
Lower Layer: IPv6     Single Hop
Session State : Up
Local Discriminator : 1 Remote Discriminator: 163
Local Address : 2001::1222/128    Remote Address: 2001::1223/128
Local Port : 49152      Remote Port: 3784
Timers in Milliseconds
Min Tx: 1000 Min Rx: 1000 Multiplier: 4
UP Count: 1 UPTIME: 00:06:03


#show bfd session ipv6 vrf raj fe80::ba6a:97ff:fed3:26c5
fe80::ba6a:97ff:fece:3bc5 detail

Session Interface Index : 10017          Interface name :xe16
Session Index : 259
Lower Layer : IPv6                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 259                Local Address :
fe80::ba6a:97ff:fed3:26
c5/128
```

```
Remote Discriminator : 257              Remote Address :
fe80::ba6a:97ff:fece:3
bc5/128
Local Port : 49152                      Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250             Min Rx: 250             Multiplier: 3
Neg Tx: 250             Neg Rx: 250             Neg detect mult: 3
Min echo Tx: 1000       Min echo Rx: 1000       Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                                    Pkt Out : n/a
Pkts Drop : 00000000000000000000                Auth Pkts Drop :
00000000000000000000
Echo Out : 00000000000000000000                 IPv6 Echo Out :
00000000000000000000
IPv6 Pkt In : 00000000000000001671              IPv6 Pkt Out : 00000000000000001675
UP Count : 1                                    UPTIME : 00:06:05

Protocol Client Info:
OSPF6-> Client ID: 5     Flags: 4
```

Table P-9 explains the output fields.

Table P-9: show BFD session fields

Entry	Description
Session Interface Index	ID number of the Interface.
Session Index	ID number for this BFD session.
Lower Layer	The lower layer protocol on which BFD is carried (IPv4 or IPv6).
Session State	The state of the session: Init: session is initializing Up: session is up. Down: session is down AdminDown: session has been administratively shutdown indefinitely
Session Type	Whether the session is Single Hop or Multiple Hop
Local Discriminator	A unique value generated by the transmitting device used to demultiplex multiple BFD sessions between the same pair of devices.
Remote Discriminator	The discriminator received from the corresponding remote device. This field is zero if the value is unknown.

Table P-9: show BFD session fields

Entry	Description
Local Address	Local address of the transmitting device.
Remote Address	Remote address of the receiving device.
Local Port	UDP port number of the transmitting device.
Remote Port	UDP port number of the receiving device.
Min Tx	Minimum transmit interval in milliseconds.
Min Rx	Minimum receive interval in milliseconds.
Min echo Tx	Minimum transmit interval for echo packets
Min echo Rx	minimum receive interval for echo packets
Neg echo intrvl	The negotiated
Storage type	Indicates the storage type: 1 = other 2 = volatile 3 = nonvolatile 4 = permanent 5 = read only
Sess down time	How long the session was down
Sess Down Reason	Reason for Session being down
Bfd GTSM	Whether Generalized TTL Security Mechanism (GTSM) is enabled or disabled on the connection.
Multiplier	The negotiated transmit interval, multiplied by this value, provides the detection time for the receiving system in asynchronous mode.
UP Count	Number of times the BFD session has been in up state.
UPTIME	Length of time this BFD session has been in the up state.
Bfd Authentication	If authentication is used, if yes, the type authentication: simple password, MD5, and so on.
Pkt In	Number of Packets received on the BFD session
Pkt Out	Number of packets transmitted on the BFD session
Pkts Drop	Number of packets dropped on the session
Auth Pkts Drop	Number of Authentication packets dropped
Echo Out	Number of Transmitted BFD Echo Packets
IPv6 Echo Out	Number of transmitted BFD Ipv6 Echo packets
IPv6 Pkt In	Number of received IPv6 BFD packets
IPv6 Pkt Out	Number of transmitted IPV6 BFD packets
Protocol Client Info	Client ID and flags

show debugging bfd

Use this command to display debugging information for BFD processes.

Command Syntax

```
show debugging bfd
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below displays the command syntax and sample output from the command.

```
#show debugging bfd
BFD debugging status:
BFD events debugging is on
BFD packet debugging is on
BFD ipc-error debugging is on
BFD ipc-event debugging is on
BFD session debugging is on
BFD nsm debugging is on
#
```

snmp restart bfd

Use this command to restart SNMP in Bidirectional Forwarding (BFD)

Command Syntax

```
snmp restart bfd
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart bfd
```

CHAPTER 2 Protocol Commands for BFD

The chapter describes the commands used to manage BFD functionality for OSPF, IS-IS and BGP.

- area virtual-link
- bfd all-interfaces
- debug bgp bfd
- debug isis bfd
- debug ospf bfd
- ip ospf bfd
- isis bfd

© 2025 IP Infusion Inc. Proprietary

area virtual-link

Use this command to enable the BFD option for a specified virtual-link neighbor.

Use the `no` form of the command to disable BFD on a virtual-link neighbor.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {fall-over bfd}
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {fall-over bfd}
```

Parameters

`A.B.C.D`	Indicate an area IP address
`<0-429467295>`	Indicate an area ID in integer format
`virtual-link`	Indicate a virtual link and its parameters
`A.B.C.D`	Indicate the IP address of the virtual link
`fall-over`	Indicate fall-over detection
`bfd`	Specify the Bidirectional Forwarding Detection (BFD)

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ospf
(config-router)#area 1 virtual-link 192.168.0.1 fall-over bfd
```

bfd all-interfaces

Use this command to enable BFD for all neighbors maintained by an OSPF process or an ISIS instance.

Use the `no` form of the command to disable BFD.

Note: This command does not apply BFD to virtual-link neighbors.

Command Syntax

```
bfd all-interfaces
no bfd all-interfaces
```

Parameters

None

Default

By default, bfd all interface is disabled

Command Mode

Configure router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ospf
(config-router)#bfd all-interfaces

#configure terminal
(config)#router isis
(config-router)#bfd all-interfaces
```

debug bgp bfd

Use this command to debug BFD processes in BGP.

Use the `no` form of the command to stop debugging.

Command Syntax

```
debug bgp bfd
no debug bgp bfd
```

Parameters

None

Command Mode

Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug bgp bfd
```

debug isis bfd

Use this command to debug BFD processes in IS-IS.

Use the `no` form of the command to stop debugging.

Command Syntax

```
debug isis bfd
no debug isis bfd
```

Parameters

None

Command Mode

Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug isis bfd
```

debug ospf bfd

Use this command to debug BFD processes in OSPF.

Use the `no` form of the command to stop debugging.

Command Syntax

```
debug ospf bfd
no debug ospf bfd
```

Parameters

None

Command Mode

Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ospf bfd
```

ip ospf bfd

Use this command to enable the BFD option for OSPF neighbors on an interface. Use the `no` form of the command to disable the BFD option for OSPF neighbors on an interface.

Note: When BFD monitoring is enabled for ospf session, protocol admin events like clear/ shutdown will cause BFD session to go to admin down.Due to this, neigbourship/adjacency down detection on peer will be according to the protocol configured dead interval and is not based on BFD interval.

Command Syntax

```
ip ospf bfd (disable|)
no ip ospf bfd (disable|)
```

Parameter

`disable`	Disable the BFD option for neighbors on an interface

Default

No default value is specified

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ip ospf bfd
```

isis bfd

Use this command to enable the BFD option for IS-IS neighbors on an interface. Use the `no` form of the command to disable the BFD option for neighbors on an interface.

Note: When BFD monitoring is enabled for ISIS session, protocol admin events like clear/ shutdown will cause BFD session to go to admin down. Due to this, neigbourship/adjacency down detection on peer will be according to the protocol configured hello hold interval and is not based on BFD interval.

Command Syntax

```
isis bfd (disable|)
no isis bfd (disable|)
```

Parameter

`disable`	Used to disable the BFD option for neighbors on an interface

Default

No default value is specified

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#isis bfd disable
```

CHAPTER 3 BFD Static Route Commands

Bidirectional Forwarding Detection (BFD) support for static routes can be configured on a static route basis, interface basis, or on a global level:

- When BFD is configured for an IPv4 or IPv6 static route, BFD provides the next-hop reachability detection for the given static route.
- When BFD is configured for an interface, BFD provides the data plane next-hop reachability information for any IPv4 or IPv6 static route configured through the given interface.
- When BFD is configured globally, BFD is applied on all interfaces with a single command. In all these cases, the BFD session update for NSM governs the state of the static routes.

This chapter includes the following commands:

- ip bfd static all-interfaces
- ip static fall-over-bfd
- ip static bfd
- ipv6 bfd static all-interfaces
- ipv6 static fall-over-bfd
- ipv6 static bfd

© 2025 IP Infusion Inc. Proprietary

ip bfd static all-interfaces

Use this command to enable BFD support for IPv4 static routes configured on all interfaces.

Use the `no` option with this command to disable BFD support for IPv4 static routes configured on all interfaces.

Command Syntax

```
ip bfd static all-interfaces
ip bfd vrf NAME static all-interfaces
no ip bfd static all-interfaces
no ip bfd vrf NAME static all-interfaces
```

Parameters

`NAME`	Enable/disable IPv4 static BFD on all interfaces bound to this user-defined Virtual Routing and Forwarding instance name.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 1.3.5.

Example

```
#configure terminal
(config)#ip bfd static all-interfaces

#configure terminal
(config)#ip  bfd  vrf vrf1 static all-interfaces
```

ip static fall-over-bfd

Use this command to enable BFD support for a specific IPv4 static route.

Use the `no` form of the command to disable the BFD support for a specific IPv4 static route.

Command Syntax

```
ip static A.B.C.D/M A.B.C.D fall-over-bfd (disable|enable)
no ip static A.B.C.D/M A.B.C.D
ip static vrf NAME A.B.C.D/M A.B.C.D fall-over-bfd (disable|enable)
no ip static vrf NAME A.B.C.D/M A.B.C.D
```

Parameters

`A.B.C.D/M`	The IPv4 destination prefix and mask length.
`A.B.C.D`	The IPv4 gateway address.
`disable`	Disable BFD.
`enable`	Enable BFD.
`NAME`	Enable/disable BFD for the IPv4 routes for this user-defined Virtual Routing and Forwarding instance name.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 1.3.5.

Example

```
#configure terminal
(config)#ip static A.B.C.D/M A.B.C.D fall-over-bfd disable

#configure terminal
(config)# ip static vrf vrf1 A.B.C.D/M A.B.C.D fall-over-bfd enable
```

ip static bfd

Use this command to enable or disable BFD support for IPv4 static route(s) configured on an interface.

Use the `no` option with this command to reset BFD support for IPv4 static route(s) configured on an interface.

Command Syntax

```
ip static bfd (disable|enable)
no ip static bfd
```

Parameters

`disable`	Disable BFD
`enable`	Enable BFD

Command Mode

Interface mode

Default

By default, BFD static route support is disabled at all levels.

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ip static bfd disable

(config)#interface eth1
(config-if)#ip static bfd enable
```

ipv6 bfd static all-interfaces

Use this command to enable BFD support for IPv6 static routes on all interfaces.

Use the `no` option with this command to disable BFD support for IPv6 static routes on all interfaces.

Command Syntax

```
ipv6 bfd static all-interfaces
ipv6 bfd vrf NAME static all-interfaces
no ipv6 bfd static all-interfaces
no ipv6 bfd vrf NAME static all-interfaces
```

Parameters

`NAME`	Enable/disable IPv6 static BFD on all interfaces bound to this user-defined Virtual Routing and Forwarding instance name.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 1.3.5.

Example

```
#configure terminal
(config)#ipv6 bfd static all-interfaces

#configure terminal
(config)#ipv6 bfd vrf vrf1 static all-interfaces
```

ipv6 static fall-over-bfd

Use this command to enable BFD support for a specific IPv6 static route.

Use the `no` option with this command to disable BFD support for a specific IPv6 static route.

Command Syntax

```
ipv6 static X:X::X:X/M X:X::X:X fall-over-bfd (disable|enable)
no ipv6 static X:X::X:X/M X:X::X:X
ipv6 static vrf NAME X:X::X:X/M X:X::X:X fall-over-bfd (disable|enable)
no ipv6 static vrf NAME X:X::X:X/M X:X::X:X
```

Parameters

`X:X::X:X/M`	The IPv6 destination prefix and mask length.
`X:X::X:X`	The IPv6 gateway address.
`NAME`	Enable/disable BFD for the IPv6 routes for this user-defined Virtual Routing and Forwarding instance name.
`disable`	Disable BFD.
`enable`	Enable BFD.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 1.3.5.

Examples

```
#configure terminal
(config)#ipv6 static X:X::X:X/M X:X::X:X fall-over-bfd disable

#configure terminal
(config)#ipv6 static X:X::X:X/M X:X::X:X fall-over-bfd enable

#configure terminal
(config)#ipv6 static vrf vrf1 X:X::X:X/M X:X::X:X fall-over-bfd disable
```

ipv6 static bfd

Use this command to disable BFD support for IPv6 static route(s) configured on an interface.

Use the `no` option with this command to reset BFD support for IPv6 static route(s) configured on an interface.

Command Syntax

```
ipv6 static bfd (disable|enable)
no ipv6 static bfd
```

Parameters

`disable`	Disable BFD
`enable`	Enable BFD

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ipv6 static bfd disable
```

Intermediate System to Intermediate System Configuration

CHAPTER 1 IS-IS IPv4

This chapter contains basic IS-IS (Intermediate System to Intermediate System) configuration examples.

Enable IS-IS on an Interface

This example shows the minimum configuration required for enabling IS-IS on an interface. R1 and R2 are two routers in the `ABC` instance connecting to the network 10.10.10.0/24. After enabling IS-IS on an interface, create a routing instance, and specify the Network Entity Title (NET). IS-IS explicitly specifies a NET to begin routing. NET is comprised of the area address and the system ID of the router.

Topology

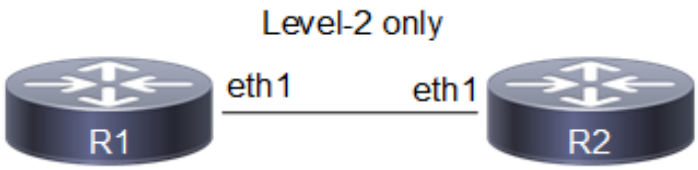

Figure 1-2: Basic IS-IS Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config)#ip address 21.21.21.2/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area `49 (ABC)`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (`ABC`).
`(config-router)#is-type level-2-only`	Configure instance as `level-2-only` routing.
`(config-router)#net 49.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config)#ip address 21.21.21.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area `49 (ABC)`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area `49 (ABC)`.

`(config-router)#is-type level-2-only`	Configure instance as `level-2-only` routing.
`(config-router)#net 49.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     24        L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     6         L2   IS-IS

R1#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0002 eth1        Up     L2   64        0000.0000.0001.01

R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0001 eth1        Up     L2   64        0000.0000.0001.01

R1#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.00dc.0b76
    IP interface address:
      21.21.21.2/24
    IPv6 interface address:
      fe80::5054:ff:fedc:b76/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 0 milliseconds

R2#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
```

```
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.002a.230a
    IP interface address:
      21.21.21.1/24
    IPv6 interface address:
      fe80::5054:ff:fe2a:230a/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 1 seconds

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    21.21.21.0/24       10          --                  eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    21.21.21.0/24       10          --                  eth1            0

R1#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0002       10         0000.0000.0002       eth1
5254.002a.230a

R2#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth1
5254.00dc.0b76
0000.0000.0002       --

R1#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000009   0x6C2D        980               0/0/0
```

```
0000.0000.0001.01-00* 0x00000003   0x1DBB         980               0/0/0
0000.0000.0002.00-00  0x0000000A   0x5444         980               0/0/0


R2#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000009   0x6C2D        942               0/0/0
0000.0000.0001.01-00  0x00000003   0x1DBB        942               0/0/0
0000.0000.0002.00-00* 0x0000000A   0x5444        944               0/0/0
```

Set Priority

This example describes how to set the priority for an interface. Set a high priority for a router to make it the Designated IS (DIS). Router R2 is configured to have a priority of 125, this is higher than the default priority (64) of R1. This makes R2 the DIS..

Topology

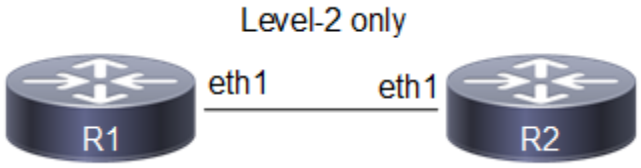

Figure 1-3: Set IS-IS Priority

Configuration

R1

`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area `49 (ABC)`.
`(config)#ip address 21.21.21.2/24`	Configure IP address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area `49` (`ABC`).
`(config-router)#is-type level-2-only`	Configure instance as `level-2-only` routing.
`(config-router)#net 49.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`(config)#interface eth1`	Enter interface mode.
`(config)#ip address 21.21.21.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area `49 (ABC)`.

`(config-if)#isis priority 125`	Specify the router priority to a higher priority (125) to make R2 the designated IS (DIS).
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area `49` (`ABC`).
`(config-router)#is-type level-2-only`	Configure instance as `level-2-only` routing.
`(config-router)#net 49.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#show clns neighbors
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     6         L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     21        L2   IS-IS

R1#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0002 eth1        Up     L2   125       0000.0000.0002.01

R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0001 eth1        Up     L2   64        0000.0000.0002.01

R1#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.00dc.0b76
    IP interface address:
      21.21.21.2/24
    IPv6 interface address:
      fe80::5054:ff:fedc:b76/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0002.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
```

```
    Next IS-IS LAN Level-2 Hello in 1 seconds

R2#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.002a.230a
    IP interface address:
      21.21.21.1/24
    IPv6 interface address:
      fe80::5054:ff:fe2a:230a/64
    Level-2 Metric: 10/10, Priority: 125, Circuit ID: 0000.0000.0002.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 737 milliseconds
```

Dynamic hostname

This example shows how to configure Dynamic Hostname for an ISIS instance. Dynamic hostname is the method of mapping name-to-systemID. It allows the routing protocol to advertise symbolic names in the IS-IS PDUs. This is done by the addition of a new TLV which allows the IS-IS routers to include the name-to-systemID mapping data in their LSPs. This allows for simple and reliable transport of name mapping across IS-IS networks. Dynamic hostname can be either the hostname of the node or the tag of the configured ISIS instance.

Note: Dynamic-hostname has to be configured on all nodes for it to take effect.

Topology

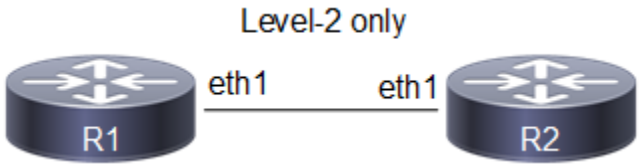

Figure 1-4: Basic dynamic hostname topology

Configuration

R1

(config)#interface eth1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config)#ip address 21.21.21.2/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0000.0000.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.

R2

(config)#interface eth1	Enter interface mode.
(config)#ip address 21.21.21.1/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R2             eth1        5254.002a.230a  Up     28        L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             eth1        5254.00dc.0b76  Up     7         L2   IS-IS

R1#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
```

```
R2              eth1        Up     L2   64         0000.0000.0001.01

R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R1             eth1        Up     L2   64         0000.0000.0001.01

R1#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       eth1        5254.002a.230a

R2#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1       eth1        5254.00dc.0b76
R2                   --

R1#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x0000000B   0x1D6B        1170              0/0/0
R1.01-00            * 0x00000004   0x1BBC        538               0/0/0
R2.00-00              0x0000000C   0x0D79        1166              0/0/0

R2#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x0000000B   0x1D6B        1078              0/0/0
R1.01-00              0x00000004   0x1BBC        445               0/0/0
R2.00-00            * 0x0000000C   0x0D79        1075              0/0/0
```

Redistribute Routes into IS-IS

In this example, the configuration causes OSPF routes to be imported into the IS-IS routing table, and advertised into the `ABC` instance.

Topology

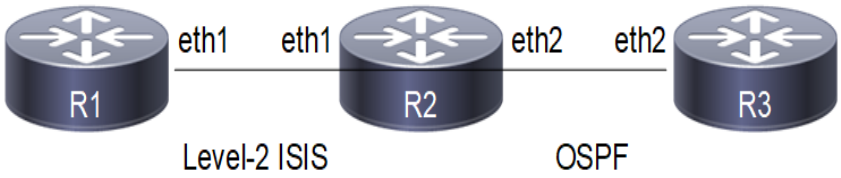

Figure 1-5: Redistribute Routes Into IS-IS

Configuration

R1

(config)#interface eth1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config)#ip address 21.21.21.2/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0000.0000.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.

R2

(config)#interface eth1	Enter interface mode.
(config)#ip address 21.21.21.1/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config)#ip address 31.31.31.1/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#redistribute ospf	Specify redistributing routes from other routing protocol (OSPF) into IS-IS.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.
(config-router)#exit	Exit interface mode.
(config)#interface lo	Configure interface lo
(config-if)#ip address 2.2.2.2/32 secondary	Configure secondary IP address to loopback interface
(config-if)#exit	Exit interface mode.

© 2025 IP Infusion Inc. Proprietary

(config)#router ospf 100	Configure OSPF routing process and specify the tag (100) which uniquely identifies the routing process
(config-router)#ospf router-id 2.2.2.2	Specify a Router ID (2.2.2.2) for the OSPF routing process.
(config-router)#network 2.2.2.2/32 area 0.0.0.0	Advertising 2.2.2.2 network
(config-router)#network 31.31.31.0/24 area 0.0.0.0	Advertising 31 network
(config-router)#exit	Exit router mode.

R3

(config)#interface eth2	Enter interface mode.
(config)#ip address 31.31.31.2/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#interface lo	Configure interface lo
(config-if)#ip address 3.3.3.3/32 secondary	Configure secondary IP address to loopback interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 100	Configure OSPF routing process and specify the tag (100) which uniquely identifies the routing process
(config-router)#ospf router-id 3.3.3.3	Specify a Router ID (3.3.3.3) for the OSPF routing process.
(config-router)#network 3.3.3.3/32 area 0.0.0.0	Advertising 3.3.3.3 network
(config-router)#network 31.31.31.0/24 area 0.0.0.0	Advertising 31 network

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R2             eth1        5254.002a.230a      Up     25        L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             eth1        5254.00dc.0b76      Up     6         L2   IS-IS

R1#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R2             eth1        Up     L2   64        0000.0000.0001.01
```

```
R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R1             eth1        Up     L2   64        0000.0000.0001.01

R1#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       eth1        5254.002a.230a

R2#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1       eth1        5254.00dc.0b76
R2                   --

R1#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000003   0x2D63        1096              0/0/0
R1.01-00            * 0x00000002   0x1FBA        1096              0/0/0
R2.00-00              0x00000004   0xEF02        1108              0/0/0

R2#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000003   0x2D63        1021              0/0/0
R1.01-00              0x00000002   0x1FBA        1021              0/0/0
R2.00-00            * 0x00000004   0xEF02        1035              0/0/0


R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   2.2.2.2/32          10          21.21.21.1          eth1            0
L2   3.3.3.3/32          10          21.21.21.1          eth1            0
C    21.21.21.0/24       10          --                  eth1            0
L2   31.31.31.0/24       10          21.21.21.1          eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid
```

```
Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
E    2.2.2.2/32          0           --                  --             0
E    3.3.3.3/32          0           --                  --             0
C    21.21.21.0/24       10          --                  eth1            0
E    31.31.31.0/24       0           --                  --             0

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L2         2.2.2.2/32 [115/10] via 21.21.21.1, eth1, 00:16:54
i L2         3.3.3.3/32 [115/10] via 21.21.21.1, eth1, 00:16:43
C            10.12.30.0/24 is directly connected, eth0, 00:24:28
C            21.21.21.0/24 is directly connected, eth1, 00:18:37
i L2         31.31.31.0/24 [115/10] via 21.21.21.1, eth1, 00:16:54
C            127.0.0.0/8 is directly connected, lo, 00:24:28

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            2.2.2.2/32 is directly connected, lo, 00:21:31
O            3.3.3.3/32 [110/2] via 31.31.31.2, eth2, 00:20:14
C            10.12.30.0/24 is directly connected, eth0, 00:27:36
C            21.21.21.0/24 is directly connected, eth1, 00:21:31
C            31.31.31.0/24 is directly connected, eth2, 00:21:31
C            127.0.0.0/8 is directly connected, lo, 00:27:36

Gateway of last resort is not set

R2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address         Interface
Instance ID
3.3.3.3          1   Full/Backup     00:00:35   31.31.31.2     eth2
0
```

```
R3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/DR         00:00:32    31.31.31.1      eth2
0

R2#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting

C  2.2.2.2/32 [1] is directly connected, lo, Area 0.0.0.0
O  3.3.3.3/32 [2] via 31.31.31.2, eth2, Area 0.0.0.0
C  31.31.31.0/24 [1] is directly connected, eth2, Area 0.0.0.0

R3#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting

O  2.2.2.2/32 [2] via 31.31.31.1, eth2, Area 0.0.0.0
C  3.3.3.3/32 [1] is directly connected, lo, Area 0.0.0.0
C  31.31.31.0/24 [1] is directly connected, eth2, Area 0.0.0.0

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            2.2.2.2/32 [110/2] via 31.31.31.1, eth2, 00:19:47
C            3.3.3.3/32 is directly connected, lo, 00:20:40
C            10.12.30.0/24 is directly connected, eth0, 00:26:28
C            31.31.31.0/24 is directly connected, eth2, 00:20:40
C            127.0.0.0/8 is directly connected, lo, 00:26:28

Gateway of last resort is not set
```

Metric

You can make a route the preferred route by changing its metric. In this example, the cost has been configured to make R3 the next hop for R1.

The default metric for each interface is 10. Interface eth3 on R2 has a metric of 20, and Interface eth2 on R3 has a metric of 30. The total cost to reach 10.10.14.0/24 (R4) through R2 and R3 is computed as follows:

R2: 10+20 = 30

R3: 10+30 = 40

In this topology, R1 chooses R2 as its next hop for destination 10.10.14.0/24.

Note: Below configuration is applicable for narrow (non-wide) metric-style. Wide metric can be configured by using the CLI's "metric-style wide" under isis instance and "isis wide-metric < 1-16777214>" under interface mode.

Topology

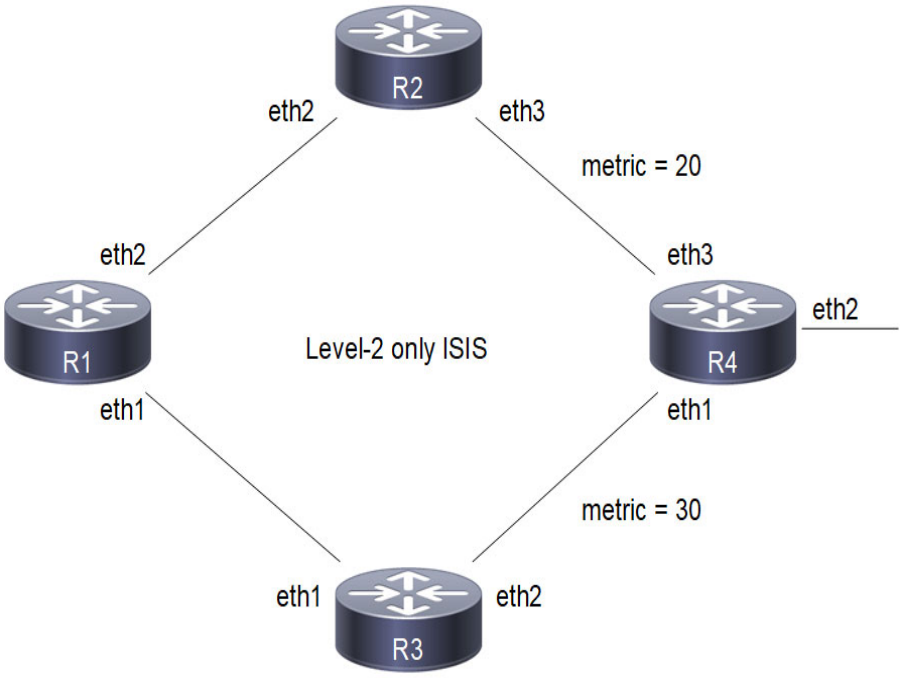

Figure 1-6: Configure IS-IS Metric

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).
`(config-if)#exit`	Exit interface mode.

(config)#interface eth2	Enter interface mode.
(config-if)#ip address 30.30.30.1/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.
(config-router)#net 49.0000.0000.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

(config)#interface eth2	Enter interface mode.
(config-if)#ip address 30.30.30.2/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)#ip address 40.40.40.1/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis metric 20	Set the value of IS-IS metric (on eth3) to 20.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.
(config-router)#net 49.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R3

(config)#interface eth1	Enter interface mode.
(config-if)#ip address 20.20.20.2/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#ip address 50.50.50.1/24	Configure IP address on interface.
(config-if)#isis metric 30	Set the value of IS-IS metric (on eth2) to 30.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).

© 2025 IP Infusion Inc. Proprietary

(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.
(config-router)#net 49.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R4

(config)#interface eth1	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#ip address 50.50.50.2/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)#ip router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#ip address 40.40.40.2/24	Configure IP address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#dynamic-hostname	Configure the hostname to be advertised for an ISIS instance.
(config-router)#net 49.0000.0000.0004.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R3             eth1        5254.00dc.2f11  Up     5         L2   IS-IS
R2             eth2        5254.007e.5ade  Up     20        L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             eth2        5254.00a1.6afe  Up     7         L2   IS-IS
R4             eth3        5254.00b1.d6fb  Up     8         L2   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
```

```
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             eth1        5254.00dc.0b76  Up     20        L2   IS-IS
R4             eth2        5254.00f5.35a4  Up     8         L2   IS-IS

R4#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R3             eth1        5254.00a8.940d  Up     25        L2   IS-IS
R2             eth3        5254.0049.c509  Up     25        L2   IS-IS

R1#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       eth2        5254.007e.5ade
R3                   10         R3       eth1        5254.00dc.2f11
R4                   30         R2       eth2        5254.007e.5ade

R2#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1       eth2        5254.00a1.6afe
R2                   --
R3                   20         R1       eth2        5254.00a1.6afe
R4                   20         R4       eth3        5254.00b1.d6fb

R3#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1       eth1        5254.00dc.0b76
R2                   20         R1       eth1        5254.00dc.0b76
R3                   --
R4                   30         R4       eth2        5254.00f5.35a4

R4#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       eth3        5254.0049.c509
                                R3       eth1        5254.00a8.940d
R2                   10         R2       eth3        5254.0049.c509
R3                   10         R3       eth1        5254.00a8.940d
R4                   --

R1#show ip isis route
```

```
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
C    30.30.30.0/24       10          --                  eth2            0
L2   40.40.40.0/24       30          30.30.30.2          eth2            0
L2   50.50.50.0/24       40          30.30.30.2          eth2            0
                                     20.20.20.2          eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   20.20.20.0/24       20          30.30.30.1          eth2            0
C    30.30.30.0/24       10          --                  eth2            0
C    40.40.40.0/24       20          --                  eth3            0
L2   50.50.50.0/24       30          40.40.40.2          eth3            0


R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
L2   30.30.30.0/24       20          20.20.20.1          eth1            0
L2   40.40.40.0/24       40          20.20.20.1          eth1            0
                                     50.50.50.2          eth2            0
C    50.50.50.0/24       30          --                  eth2            0

R4#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   20.20.20.0/24       20          50.50.50.1          eth1            0
L2   30.30.30.0/24       20          40.40.40.1          eth3            0
C    40.40.40.0/24       10          --                  eth3            0
C    50.50.50.0/24       10          --                  eth1            0

R1#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
```

```
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.00dc.0b76
    IP interface address:
      20.20.20.1/24
    IPv6 interface address:
      fe80::5054:ff:fedc:b76/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0003.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 5 seconds
eth2 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x02
    Extended Local circuit ID: 0x00000004
    Local SNPA: 5254.00a1.6afe
    IP interface address:
      30.30.30.1/24
    IPv6 interface address:
      fe80::5054:ff:fea1:6afe/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.02
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 183 milliseconds

R2#show isis interface
eth2 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000004
    Local SNPA: 5254.007e.5ade
    IP interface address:
      30.30.30.2/24
    IPv6 interface address:
      fe80::5054:ff:fe7e:5ade/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.02
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 706 milliseconds
eth3 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x02
    Extended Local circuit ID: 0x00000005
    Local SNPA: 5254.0049.c509
    IP interface address:
      40.40.40.1/24
    IPv6 interface address:
      fe80::5054:ff:fe49:c509/64
    Level-2 Metric: 20/10, Priority: 64, Circuit ID: 0000.0000.0004.02
```

```
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 2 seconds

R3#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.00dc.2f11
    IP interface address:
      20.20.20.2/24
    IPv6 interface address:
      fe80::5054:ff:fedc:2f11/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0003.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 2 seconds
eth2 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x02
    Extended Local circuit ID: 0x00000004
    Local SNPA: 5254.00a8.940d
    IP interface address:
      50.50.50.1/24
    IPv6 interface address:
      fe80::5054:ff:fea8:940d/64
    Level-2 Metric: 30/10, Priority: 64, Circuit ID: 0000.0000.0004.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 3 seconds

R4#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00000003
    Local SNPA: 5254.00f5.35a4
    IP interface address:
      50.50.50.2/24
    IPv6 interface address:
      fe80::5054:ff:fef5:35a4/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0004.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 0 milliseconds
eth3 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x02
```

```
Extended Local circuit ID: 0x00000005
Local SNPA: 5254.00b1.d6fb
IP interface address:
  40.40.40.2/24
IPv6 interface address:
  fe80::5054:ff:feb1:d6fb/64
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0004.02
Number of active level-2 adjacencies: 1
Level-2 LSP MTU: 1492
Next IS-IS LAN Level-2 Hello in 0 milliseconds
```

L1-L2 Area Routing with a Single Instance

IS-IS supports a two-level hierarchy for handling and scaling the functionality of large networks. The Level-1 (L1) area is mainly for Leaf networks, and the Level-2 (L2) area is the backbone area connecting Level-1 areas. In this example, R3 and R4 are configured as Level-1 routers, and reside in the Level-1 area. R1 and R2 are configured as Level-1-2 routers, and connect these two Level-1 areas with a backbone Level-2 area. You can configure Level-1-2 routers with single or multiple instances: This configuration shows the single-instance version of the Level-1-2 router.

Topology

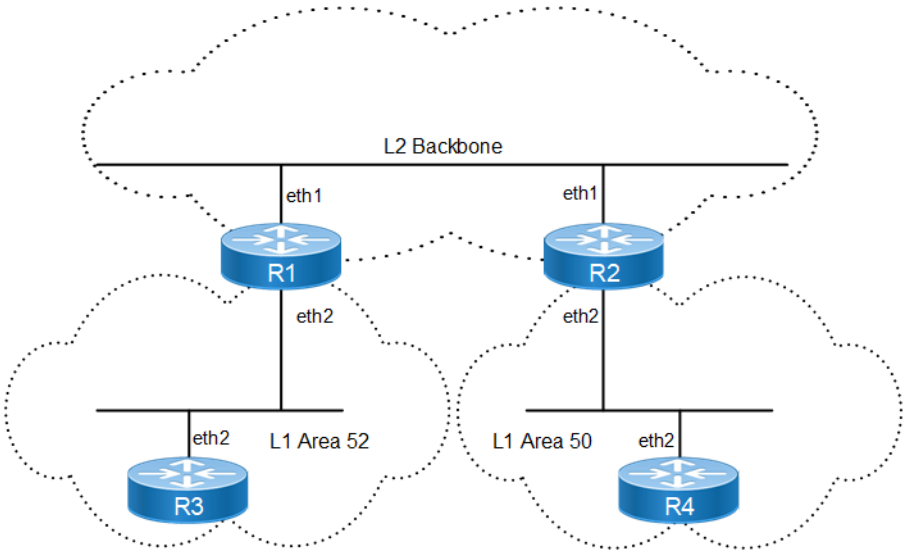

Figure 1-7: Single-Instance L1-L2 Area Routing

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on the interface `eth1` for area `ABC`.

`#configure terminal`	Enter configure mode.
`(config-if)#isis circuit-type level-2-only`	Set the circuit type for the interface `eth1`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 30.30.30.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ABC`	Enable IS-IS routing on the interface `eth2` for area `ABC`.
`(config-if)#isis circuit-type level-1`	Set the circuit type for interface `eth2` to level 1.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area `ABC`.
`(config-router)#net 52.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip router isis bb`	Enable IS-IS routing on the interface `eth1` for area `bb`.
`(config-if)#ip address 20.20.20.2/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-2-only`	Set the circuit type for the interface `eth1` to level-2 only.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 40.40.40.1/24`	Configure IP address on interface.
`(config-if)#ip router isis bb`	Enable IS-IS routing on interface `eth2` for area `bb`.
`(config-if)#isis circuit-type level-1`	Set the circuit type for interface `eth2` to level 1.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis bb`	Create an IS-IS routing instance for area `bb`.
`(config-router)#net 50.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R3

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 30.30.30.2/24`	Configure IP address on interface.
`(config-if)#ip router isis xyz`	Enable IS-IS routing on the interface `eth2` for area `xyz`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis xyz`	Create an IS-IS routing instance for area `xyz`.
`(config-router)#is-type level-1`	Set the IS level for this area (`xyz`) as `level-1`.
`(config-router)#net 52.0000.0000.0003.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R4

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 40.40.40.2/24`	Configure IP address on interface.

`(config-if)#ip router isis aa`	Enable IS-IS routing on the interface `eth2` for area `aa`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis aa`	Create an IS-IS routing instance for area `aa`.
`(config-router)#is-type level-1`	Set the IS level for this area (`aa`) as `level-1`.
`(config-router)#net 50.0000.0000.0004.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     20        L2   IS-IS
0000.0000.0003 eth2        5254.00a8.940d  Up     6         L1   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag bb:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     8         L2   IS-IS
0000.0000.0004 eth2        5254.00e2.aece  Up     7         L1   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag xyz:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth2        5254.00a1.6afe  Up     23        L1   IS-IS

R4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag aa:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.007e.5ade  Up     25        L1   IS-IS

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid
```

```
Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
C    30.30.30.0/24       10          --                  eth2            0
L2   40.40.40.0/24       20          20.20.20.2          eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag bb:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
L2   30.30.30.0/24       20          20.20.20.1          eth1            0
C    40.40.40.0/24       10          --                  eth2            0

R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag xyz:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L1   0.0.0.0/0           10          30.30.30.1          eth2            0
C    30.30.30.0/24       10          --                  eth2            0

R4#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag aa:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L1   0.0.0.0/0           10          40.40.40.1          eth2            0
C    40.40.40.0/24       10          --                  eth2            0

R1#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0003       10         0000.0000.0003       eth2
5254.00a8.940d

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0002       10         0000.0000.0002       eth1
5254.002a.230a

R2#show isis topology
```

```
Tag bb:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0002       --
0000.0000.0004       10         0000.0000.0004       eth2
5254.00e2.aece

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth1
5254.00dc.0b76
0000.0000.0002       --

R3#show isis topology

Tag xyz:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth2
5254.00a1.6afe
0000.0000.0003       --

R4#show isis topology

Tag aa:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0002       10         0000.0000.0002       eth2
5254.007e.5ade
0000.0000.0004       --

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.12.30.0/24 is directly connected, eth0, 00:27:08
C            20.20.20.0/24 is directly connected, eth1, 00:16:57
C            30.30.30.0/24 is directly connected, eth2, 00:15:48
i L2         40.40.40.0/24 [115/20] via 20.20.20.2, eth1, 00:15:05
C            127.0.0.0/8 is directly connected, lo, 00:27:08

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
```

```
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.12.30.0/24 is directly connected, eth0, 00:27:17
C            20.20.20.0/24 is directly connected, eth1, 00:17:13
i L2         30.30.30.0/24 [115/20] via 20.20.20.1, eth1, 00:16:18
C            40.40.40.0/24 is directly connected, eth2, 00:15:36
C            127.0.0.0/8 is directly connected, lo, 00:27:17

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 30.30.30.1 to network 0.0.0.0

i*L1         0.0.0.0/0 [115/10] via 30.30.30.1, eth2, 00:16:47
C            10.12.30.0/24 is directly connected, eth0, 00:27:46
C            30.30.30.0/24 is directly connected, eth2, 00:16:52
C            127.0.0.0/8 is directly connected, lo, 00:27:46

R4#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 40.40.40.1 to network 0.0.0.0

i*L1         0.0.0.0/0 [115/10] via 40.40.40.1, eth2, 00:16:26
C            10.12.30.0/24 is directly connected, eth0, 00:27:20
C            40.40.40.0/24 is directly connected, eth2, 00:16:36
C            127.0.0.0/8 is directly connected, lo, 00:27:20

R1#show isis database
Tag ABC:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000005   0xE66E        1165              1/0/0
0000.0000.0003.00-00  0x00000004   0xDC80        1164              0/0/0
0000.0000.0003.01-00  0x00000002   0x10C8        1163              0/0/0

IS-IS Level-2 Link State Database:
```

```
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000005   0xD0D8        1148              0/0/0
0000.0000.0001.01-00* 0x00000002   0x1FBA        1109              0/0/0
0000.0000.0002.00-00  0x00000005   0x7219        1189              0/0/0

R2#show isis database
Tag bb:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00* 0x00000005   0x9583        1179              1/0/0
0000.0000.0004.00-00  0x00000004   0x8B95        1177              0/0/0
0000.0000.0004.01-00  0x00000002   0x2FA6        1177              0/0/0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000005   0xD0D8        1116              0/0/0
0000.0000.0001.01-00  0x00000002   0x1FBA        1078              0/0/0
0000.0000.0002.00-00* 0x00000005   0x7219        1160              0/0/0

R3#show isis database
Tag xyz:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000005   0xE66E        1094              1/0/0
0000.0000.0003.00-00* 0x00000004   0xDC80        1095              0/0/0
0000.0000.0003.01-00* 0x00000002   0x10C8        1094              0/0/0

R4#show isis database
Tag aa:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00  0x00000005   0x9583        1105              1/0/0
0000.0000.0004.00-00* 0x00000004   0x8B95        1105              0/0/0
0000.0000.0004.01-00* 0x00000002   0x2FA6        1105              0/0/0
```

L1-L2 Area Routing with Multiple Instances

IS-IS supports a two-level hierarchy for handling and scaling the functionality of large networks. The Level-1 (L1) area is mainly for Leaf networks, and the Level-2 (L2) area is the backbone area connecting Level-1 areas. In this example, R3 and R4 are configured as Level-1 routers, and reside in the Level-1 area. R1 and R2 are configured as Level-1-2 routers, and connect these two Level-1 areas with a backbone Level-2 area. You can configure Level-1-2 routers with single or multiple instances: This configuration shows the multiple-instance version of the Level-1-2 router.

Topology

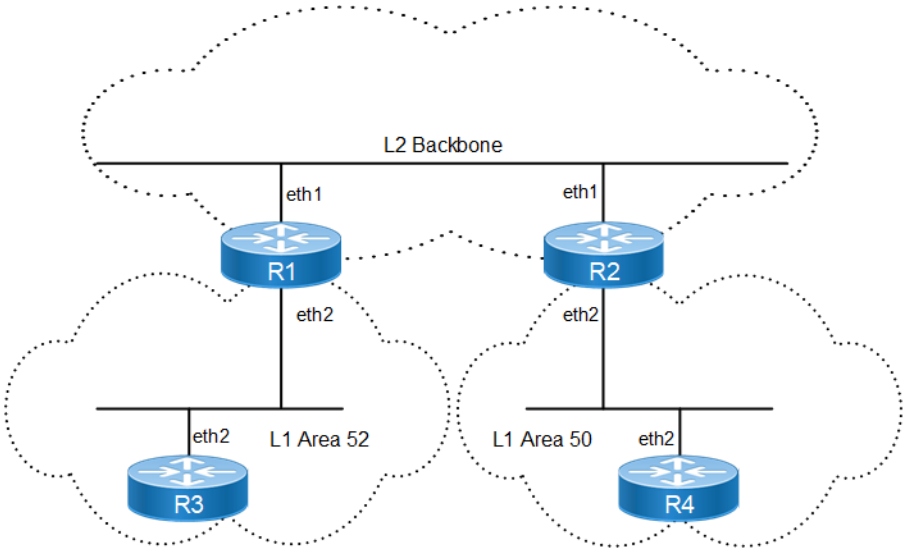

Figure 1-8: Multiple-Instance L1-L2 Area Routing

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#ip router isis aaa`	Enable IS-IS routing on interface `eth1` for area `aaa`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis aaa`	Create an IS-IS routing instance for area `aaa`.
`(config-router)#is-type level-2-only`	Set the IS level for this area (`aaa`) as `level-2-only`.
`(config-router)#net bb.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 30.30.30.1/24`	Configure IP address on interface.
`(config-if)#ip router isis ccc`	Enable IS-IS routing on interface `eth2` for area `ccc`.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ccc`	Create an IS-IS routing instance for area `ccc`.
`(config-router)#is-type level-1`	Set the IS level for this area (`ccc`) as `level-1`.
`(config-router)#net cc.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

(config)#interface eth1	Enter interface mode.
(config-if)#ip address 20.20.20.2/24	Configure IP address on interface.
(config-if)#ip router isis bb	Enable IS-IS routing on interface eth1 for area bb.
(config-if)#exit	Exit interface mode.
(config)#router isis bb	Create an IS-IS routing instance for area bb.
(config-router)#is-type level-2-only	Set the IS level for this area (bb) as level-2-only.
(config-router)#net bb.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ip address 40.40.40.1/24	Configure IP address on interface.
(config-if)#ip router isis ABC	Enable IS-IS routing on interface eth2 for area ABC.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area ABC.
(config-router)#is-type level-1	Set the IS level for this area (ABC) as level-1.
(config-router)#net cc.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R3

(config)#interface eth2	Enter interface mode.
(config-if)#ip address 30.30.30.2/24	Configure IP address on interface.
(config-if)#ip router isis xyz	Enable IS-IS routing on interface eth2 for area xyz.
(config-if)#exit	Exit interface mode.
(config)#router isis xyz	Create an IS-IS routing instance for area xyz.
(config-router)#is-type level-1	Set the IS level for this area (xyz) as level-1.
(config-router)#net 52.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R4

(config)#interface eth2	Enter interface mode.
(config-if)#ip address 40.40.40.2/24	Configure IP address on interface.
(config-if)#ip router isis aa	Enable IS-IS routing on interface eth2 for area aa.
(config-if)#exit	Exit interface mode.
(config)#router isis aa	Create an IS-IS routing instance for area aa.
(config-router)#is-type level-1	Set the IS level for this area (aa) as level-1.
(config-router)#net 52.0000.0000.0004.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

© 2025 IP Infusion Inc. Proprietary

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag aaa:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     25        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 0
Total number of adjacencies: 0
Tag ccc:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 0
Total number of adjacencies: 0
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag bb:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     6         L2   IS-IS

R1#show clns is-neighbors

Tag aaa:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0002 eth1        Up     L2   64        0000.0000.0001.01

Tag ccc:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id

R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id

Tag bb:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0001 eth1        Up     L2   64        0000.0000.0001.01

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid
```

```
Tag aaa:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0


Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ccc:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    30.30.30.0/24       10          --                  eth2            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag ABC:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    40.40.40.0/24       10          --                  eth2            0


Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag bb:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0

R1#show isis topology

Tag aaa:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0002       10         0000.0000.0002       eth1
5254.002a.230a


Tag ccc:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0002       --

R2#show isis topology

Tag ABC:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0003       --


Tag bb:  VRF : default
IS-IS paths to level-2 routers
```

```
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth1
5254.00dc.0b76
0000.0000.0002       --

R1#show isis database
Tag aaa:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000002   0x181D        1003              0/0/0
0000.0000.0001.01-00* 0x00000001   0x21B9        1003              0/0/0
0000.0000.0002.00-00  0x00000005   0x1818        1080              0/0/0

Tag ccc:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00* 0x00000001   0xDFA5        685               0/0/0

R2#show isis database
Tag ABC:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0003.00-00* 0x00000002   0xD571        696               0/0/0

Tag bb:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000002   0x181D        938               0/0/0
0000.0000.0001.01-00  0x00000001   0x21B9        938               0/0/0
0000.0000.0002.00-00* 0x00000005   0x1818        1017              0/0/0

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.12.30.0/24 is directly connected, eth0, 01:37:50
C            20.20.20.0/24 is directly connected, eth1, 00:10:13
C            30.30.30.0/24 is directly connected, eth2, 00:10:13
C            127.0.0.0/8 is directly connected, lo, 01:37:50

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
```

```
        * - candidate default

IP Route Table for VRF "default"
C            10.12.30.0/24 is directly connected, eth0, 01:37:26
C            20.20.20.0/24 is directly connected, eth1, 00:05:03
C            40.40.40.0/24 is directly connected, eth2, 00:08:52
C            127.0.0.0/8 is directly connected, lo, 01:37:26

Gateway of last resort is not set
```

Route Leaking

Route leaking is defined in RFC 2966. For Level-1 (L1) routers, only level-1 routes are populated in the routing table. The L1 router has a default route to the nearest Level-1/Level-2 (L1/L2) router: This could result in sub-optimal routing in certain scenarios. Route leaking causes an L1/L2 router to advertise the level-2 routes in its database to the L1 router, thus allowing the L1 router to acknowledge the prefixes advertised by the Level-2 (L2) router. In this way, the L1 router has the ability to learn the true cost to reach other areas.

In the following example, R1 is the L1 router, R2 is the L1/L2 router doing the route leaking, and R3 is the L2 router. The following configuration is given only for R2, assuming that the adjacency with R1 and R3 are already up, and the route tables with appropriate routes are already populated.

Topology

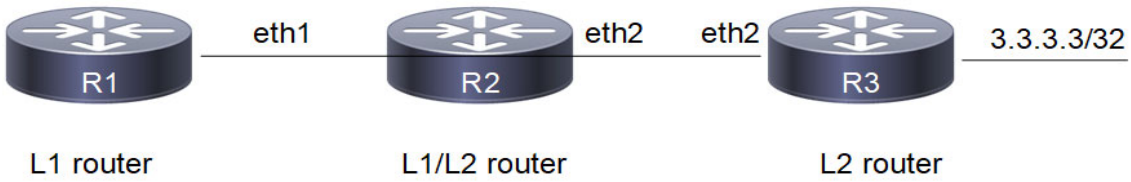

Figure 1-9: Route Leaking Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#net 49.0001.0000.0000.0001.00`	Define the NET address.
`(config-router)#is-type level-1`	Configure instance as level-1-only routing.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth1`	Specify the interface (eth1)to configure and enter Interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-1`	Set the circuit type as level-1 for the interface
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R2).

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Specify the interface (eth1)to configure and enter Interface mode.
`(config-if)#ip address 20.20.20.2/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R1).
`(config-if)#isis circuit-type level-1`	Configure instance as level-1-only routing.
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#interface eth2`	Specify the interface (eth2)to configure and enter Interface mode.
`(config-if)#ip address 30.30.30.1/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth2 (connected to R3).
`(config-if)#isis circuit-type level-2-only`	Configure instance as level-2-only routing.
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#net 49.0001.0000.0000.0002.00`	Define the NET address.
`(config-router-af)#redistribute isis level-2 into level-1`	Enable redistribution of isis routes from level-2 into level-1

R3

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the interface (lo)to configure and enter Interface mode.
`(config-if)#ip address 3.3.3.3/32 secondary`	Configure IP address on loopback interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface lo
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0001.0000.0000.0003.00`	Define the NET address.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth2`	Specify the interface (eth2)to configure and enter Interface mode.
`(config-if)#ip address 30.30.30.2/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-2-only`	Set the circuit type as level-2-only for the interface
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R2).

In the example, route, `i ia 3.3.3.3/32 [115/30] via 20.20.20.2, eth1, 00:12:29`, is the L2 route leaked by the L1/L2 router into the L1 router.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a      Up     21        L1   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76      Up     5         L1   IS-IS
0000.0000.0003 eth2        5254.00a8.940d      Up     6         L2   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.007e.5ade      Up     21        L2   IS-IS

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
ia   3.3.3.3/32          30          20.20.20.2          eth1            0
C    20.20.20.0/24       10          --                  eth1            0
ia   30.30.30.0/24       20          20.20.20.2          eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid
```

```
Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   3.3.3.3/32          20          30.30.30.2          eth2            0
C    20.20.20.0/24       10          --                  eth1            0
C    30.30.30.0/24       10          --                  eth2            0

R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    3.3.3.3/32          10          --                  lo            0
L2   20.20.20.0/24       20          30.30.30.1          eth2            0
C    30.30.30.0/24       10          --                  eth2            0

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i ia         3.3.3.3/32 [115/30] via 20.20.20.2, eth1, 00:20:53
C            10.12.30.0/24 is directly connected, eth0, 01:02:10
C            20.20.20.0/24 is directly connected, eth1, 00:48:08
i ia         30.30.30.0/24 [115/20] via 20.20.20.2, eth1, 00:23:30
C            127.0.0.0/8 is directly connected, lo, 01:02:10

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L2         3.3.3.3/32 [115/20] via 30.30.30.2, eth2, 00:21:07
```

```
C            10.12.30.0/24 is directly connected, eth0, 01:01:55
C            20.20.20.0/24 is directly connected, eth1, 00:48:12
C            30.30.30.0/24 is directly connected, eth2, 00:48:12
C            127.0.0.0/8 is directly connected, lo, 01:01:55

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            3.3.3.3/32 is directly connected, lo, 00:21:25
C            10.12.30.0/24 is directly connected, eth0, 01:01:26
i L2         20.20.20.0/24 [115/20] via 30.30.30.1, eth2, 00:24:06
C            30.30.30.0/24 is directly connected, eth2, 00:48:13
C            127.0.0.0/8 is directly connected, lo, 01:01:26

Gateway of last resort is not set

R1#show isis database
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x0000000C   0xE4B5        642               0/0/0
0000.0000.0001.01-00* 0x00000007   0x13C3        642               0/0/0
0000.0000.0002.00-00  0x00000012   0x8AC8        804               0/0/0

R2#show isis database
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000003   0xF6AC        304               0/0/0
0000.0000.0001.01-00  0x00000002   0x1DBE        304               0/0/0
0000.0000.0002.00-00* 0x00000009   0x2ECA        358               0/0/0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00* 0x00000007   0x2F5A        353               0/0/0
0000.0000.0003.00-00  0x00000003   0x25E6        347               0/0/0
0000.0000.0003.02-00  0x00000002   0x24B0        347               0/0/0

R3#show isis database
Tag 1:  VRF : default
```

```
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00  0x00000007   0x2F5A        335               0/0/0
0000.0000.0003.00-00* 0x00000003   0x25E6        331               0/0/0
0000.0000.0003.02-00* 0x00000002   0x24B0        331               0/0/0


R1#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0002       10         0000.0000.0002       eth1        5254.002a.230a

R2#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth1        5254.00dc.0b76
0000.0000.0002       --

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0002       --
0000.0000.0003       10         0000.0000.0003       eth2        5254.00a8.940d

R3#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0002       10         0000.0000.0002       eth2        5254.007e.5ade
0000.0000.0003       --
```

Route Summarization

Route summarization makes the routing table smaller, but still allows complete IP connectivity, if everything is configured properly.

The following example consists of a three-router topology, in which R2 is doing the summarization. In this example, R1 is the L1 router, R2 is the L1/L2 router doing the summarization, and R3 is the L2 router. The following configuration is given only for R2, assuming that the adjacencies with R1 and R3 are already up, and the route tables with the appropriate routes are already populated.

Topology

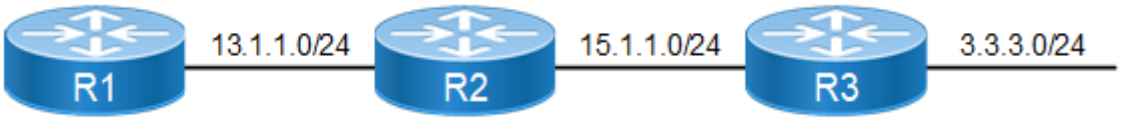

Figure 1-10: Route Summarization Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#net 49.0001.0000.0000.0001.00`	Define the NET address.
`(config-router)#is-type level-1`	Configure instance as level-1.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth1`	Specify the interface (eth1)to configure and enter Interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-1`	Set the circuit type as level-1 for the interface
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R2).

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Specify the interface (eth1)to configure and enter Interface mode.
`(config-if)#isis circuit-type level-1`	Set the circuit type as level-1 for the interface
`(config-if)#ip address 20.20.20.2/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R1).
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#interface eth2`	Specify the interface (eth2)to configure and enter Interface mode.
`(config-if)#ip address 30.30.30.1/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-2-only`	Set the circuit type as level-2-only for the interface
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth2 (connected to R3).
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#net 49.0001.0000.0000.0002.00`	Define the NET address.

© 2025 IP Infusion Inc. Proprietary

(config-router-af)#redistribute isis level-2 into level-1	Enable redistribution of isis routes from level-2 into level-1
(config-router-af)#summary-address 66.66.66.0/24 level-1 metric 50	Configure the summary address to summarize IP reachability information.

R3

#configure terminal	Enter configure mode.
(config)#ip route 66.66.66.1/32 eth2	Configure ip static route.
(config)#ip route 66.66.66.2/32 eth2	Configure ip static route.
(config)#ip route 66.66.66.3/32 eth2	Configure ip static route.
(config)#router isis 1	Create an IS-IS routing instance (1).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0001.0000.0000.0003.00	Define the NET address.
(config-router)#redistribute static	Enable redistribution of static routes into ISIS instance.
(config-router)#exit	Exit router mode.
(config)#interface eth2	Specify the interface (eth2)to configure and enter Interface mode.
(config-if)#ip address 30.30.30.2/24	Configure IP address on interface.
(config-if)#isis circuit-type level-2-only	Set the circuit type as level-2-only for the interface
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth1 (connected to R2).

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     20        L1   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     6         L1   IS-IS
0000.0000.0003 eth2        5254.00a8.940d  Up     7         L2   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 0
```

```
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.007e.5ade      Up     21        L2   IS-IS

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
ia   30.30.30.0/24       20          20.20.20.2          eth1            0
ia   66.66.66.0/24       60          20.20.20.2          eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0
C    30.30.30.0/24       10          --                  eth2            0
D    66.66.66.0/24       0           --                  --
L2   66.66.66.1/32       10          30.30.30.2          eth2            0
L2   66.66.66.2/32       10          30.30.30.2          eth2            0
L2   66.66.66.3/32       10          30.30.30.2          eth2            0

R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   20.20.20.0/24       20          30.30.30.1          eth2            0
C    30.30.30.0/24       10          --                  eth2            0
E    66.66.66.1/32       0           --                  --              0
E    66.66.66.2/32       0           --                  --              0
E    66.66.66.3/32       0           --                  --              0


R1#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000004   0xF4AD        850               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.1
  Metric:   10         IS 0000.0000.0001.01
```

```
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0001.01-00* 0x00000003   0x1BBF        850               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00  0x00000010   0xB5E0        1165              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.2
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IP 20.20.20.0 255.255.255.0
  Metric:   10         IP-Interarea 30.30.30.0 255.255.255.0
  Metric:   50         IP-External 66.66.66.0     255.255.255.0


R2#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000004   0xF4AD        820               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.1
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0001.01-00  0x00000003   0x1BBF        820               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000010   0xB5E0        1137              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.2
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IP 20.20.20.0 255.255.255.0
  Metric:   10         IP-Interarea 30.30.30.0 255.255.255.0
  Metric:   50         IP-External 66.66.66.0     255.255.255.0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00* 0x00000005   0x1577        838               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IP 30.30.30.0 255.255.255.0
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0003.00-00  0x0000000B   0xFED3        1160              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IP 30.30.30.0 255.255.255.0
  Metric:   0          IP-External 66.66.66.1     255.255.255.255
  Metric:   0          IP-External 66.66.66.2     255.255.255.255
  Metric:   0          IP-External 66.66.66.3     255.255.255.255
0000.0000.0003.01-00  0x00000003   0x29AB        837               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00
```

```
R3#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00  0x00000005   0x1577        818               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.1
  Metric:   10          IS 0000.0000.0003.01
  Metric:   10          IP 30.30.30.0 255.255.255.0
  Metric:   10          IP 20.20.20.0 255.255.255.0
0000.0000.0003.00-00* 0x0000000B   0xFED3        1142              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.2
  Metric:   10          IS 0000.0000.0003.01
  Metric:   10          IP 30.30.30.0 255.255.255.0
  Metric:   0           IP-External 66.66.66.1     255.255.255.255
  Metric:   0           IP-External 66.66.66.2     255.255.255.255
  Metric:   0           IP-External 66.66.66.3     255.255.255.255
0000.0000.0003.01-00* 0x00000003   0x29AB        819               0/0/0
  Metric:   0           IS 0000.0000.0003.00
  Metric:   0           IS 0000.0000.0002.00
```

IS-IS Distance

Administrative distance in IS-IS can be configured for a specified source ID or for all routes.

This example shows configuring the IS-IS administrative distance for the IPv4 address family.

Topology

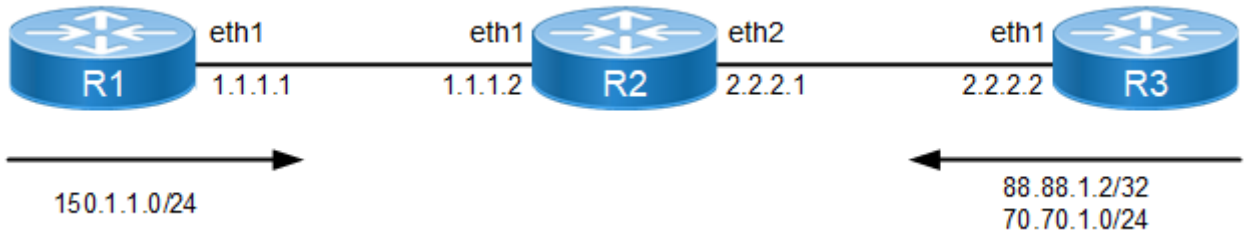

Figure 1-11: IS-IS Distance Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 1.1.1.1/24`	Assign the IP address on this interface (`eth1`).

(config-if)#ip router isis 1	Enable IS-IS routing on interface eth1.
(config-if)#exit	Exit interface mode.
(config)#ip route 150.1.1.0/24 eth1	Configure static routes.
(config)#router isis 1	Create an IS-IS routing instance (1).
(config-router)#net 49.0001.0000.0001.00	Set a Network Entity Title (NET) for this instance, specifying the address and the system ID.
(config-router)#redistribute static	Redistribute the static routes.

R2

#configure terminal	Enter configure mode
(config)#ip access-list DIST	Enter access list mode
(config-ip-acl)#permit ipip 88.88.1.2/32 any	Create an access list to permit the 88.88.1.2/32 route from R3.
(config-ip-acl)#exit	Exit access list mode
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 1.1.1.2/24	Assign the IP address on this interface (eth1).
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth1
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode.
(config-if)#ip address 2.2.2.1/24	Assign the IP address on this interface (eth2).
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth2
(config-if)#exit	Exit interface mode
(config)#router isis 1	Create an IS-IS routing instance (1).
(config-router)#net 49.0001.0000.0002.00	Specify the NET address.
(config-router)#distance 100	Configure the administrative distance for all routes received from R1 and R2.
(config-router)#distance 20 0001.0000.0001	Configure the administrative distance for all routes received from R1. This command overwrites the applied distance, 100, and will apply distance 20 for all routes received from R1.
(config-router)#distance 30 0001.0000.0003 DIST	Configure the distance, 30, to the route, 88.88.1.2/32, received from R3. All other routes from R3 (for example, 70.70.1.0/24) will have the distance applied as 100. If the distance, 100, is not configured, all other routes will have a default distance of 115.

R3

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 2.2.2.2/24	Assign the IP address on this interface (eth1).
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth1.
(config-if)#exit	Exit interface mode.
(config)#ip route 88.88.1.2/32 eth2	Configure static routes
(config)#ip route 70.70.1.0/24 eth2	Configure static routes

(config)#router isis 1	Create an IS-IS routing instance (1).
(config-router)#net 49.0001.0000.0003.00	Specify the NET address.
(config-router)#redistribute static	Redistribute the static routes.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0001.0000.0002 eth1        5254.002a.230a  Up     18        L1   IS-IS
                                           Up     18        L2   IS-IS
R2#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 2
Total number of adjacencies: 4
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0001.0000.0001 eth1        5254.00dc.0b76  Up     7         L1   IS-IS
                                           Up     7         L2   IS-IS
0001.0000.0003 eth2        5254.00a8.940d  Up     8         L1   IS-IS
                                           Up     8         L2   IS-IS
R3#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0001.0000.0002 eth2        5254.007e.5ade  Up     20        L1   IS-IS
                                           Up     20        L2   IS-IS
R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    1.1.1.0/24          10          --                  eth1            0
L1   2.2.2.0/24          20          1.1.1.2             eth1            0
L2   70.70.1.0/24        20          1.1.1.2             eth1            0
L2   88.88.1.2/32        20          1.1.1.2             eth1            0
E    150.1.1.0/24        0           --                  --             0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid
```

```
Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    1.1.1.0/24          10          --                  eth1            0
C    2.2.2.0/24          10          --                  eth2            0
L2   70.70.1.0/24        10          2.2.2.2             eth2            0
L2   88.88.1.2/32        10          2.2.2.2             eth2            0
L2   150.1.1.0/24        10          1.1.1.1             eth1            0

R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L1   1.1.1.0/24          20          2.2.2.1             eth2            0
C    2.2.2.0/24          10          --                  eth2            0
E    70.70.1.0/24        0           --                  --             0
E    88.88.1.2/32        0           --                  --             0
L2   150.1.1.0/24        20          2.2.2.1             eth2            0



R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, eth1, 00:30:56
i L1         2.2.2.0/24 [115/20] via 1.1.1.2, eth1, 00:26:01
C            10.12.30.0/24 is directly connected, eth0, 00:33:46
i L2         70.70.1.0/24 [115/20] via 1.1.1.2, eth1, 00:21:39
i L2         88.88.1.2/32 [115/20] via 1.1.1.2, eth1, 00:15:04
C            127.0.0.0/8 is directly connected, lo, 00:33:46
S            150.1.1.0/24 [1/0] is directly connected, eth1, 00:29:03

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.0/24 is directly connected, eth1, 00:26:46
```

```
C            2.2.2.0/24 is directly connected, eth2, 00:26:30
C            10.12.30.0/24 is directly connected, eth0, 00:33:21
i L2         70.70.1.0/24 [100/10] via 2.2.2.2, eth2, 00:21:55
i L2         88.88.1.2/32 [30/10] via 2.2.2.2, eth2, 00:15:09
C            127.0.0.0/8 is directly connected, lo, 00:33:21
i L2         150.1.1.0/24 [10/10] via 1.1.1.1, eth1, 00:25:53

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.0/24 [115/20] via 2.2.2.1, eth2, 00:22:56
C            2.2.2.0/24 is directly connected, eth2, 00:23:01
C            10.12.30.0/24 is directly connected, eth0, 00:33:57
S            70.70.1.0/24 [1/0] is directly connected, eth2, 00:23:01
S            88.88.1.2/32 [1/0] is directly connected, eth2, 00:16:07
C            127.0.0.0/8 is directly connected, lo, 00:33:57
i L2         150.1.1.0/24 [115/20] via 2.2.2.1, eth2, 00:22:42

Gateway of last resort is not set

R1#show isis database
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00* 0x00000003   0x2AEC        448               0/0/0
0001.0000.0001.01-00* 0x00000002   0x32A4        448               0/0/0
0001.0000.0002.00-00  0x00000004   0x5A80        698               0/0/0
0001.0000.0003.00-00  0x00000006   0xE820        702               0/0/0
0001.0000.0003.01-00  0x00000002   0x3E94        698               0/0/0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00* 0x00000008   0xB20F        703               0/0/0
0001.0000.0001.01-00* 0x00000002   0x32A4        448               0/0/0
0001.0000.0002.00-00  0x00000004   0x5A80        698               0/0/0
0001.0000.0003.00-00  0x0000000A   0xB2CE        1108              0/0/0
0001.0000.0003.01-00  0x00000002   0x3E94        698               0/0/0

R2#show isis database
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00  0x00000003   0x2AEC        402               0/0/0
0001.0000.0001.01-00  0x00000002   0x32A4        401               0/0/0
0001.0000.0002.00-00* 0x00000004   0x5A80        653               0/0/0
0001.0000.0003.00-00  0x00000006   0xE820        656               0/0/0
0001.0000.0003.01-00  0x00000002   0x3E94        652               0/0/0
```

```
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00  0x00000008   0xB20F        657               0/0/0
0001.0000.0001.01-00  0x00000002   0x32A4        401               0/0/0
0001.0000.0002.00-00* 0x00000004   0x5A80        653               0/0/0
0001.0000.0003.00-00  0x0000000A   0xB2CE        1062              0/0/0
0001.0000.0003.01-00  0x00000002   0x3E94        652               0/0/0

R3#show isis database
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00  0x00000003   0x2AEC        317               0/0/0
0001.0000.0001.01-00  0x00000002   0x32A4        317               0/0/0
0001.0000.0002.00-00  0x00000004   0x5A80        568               0/0/0
0001.0000.0003.00-00* 0x00000006   0xE820        573               0/0/0
0001.0000.0003.01-00* 0x00000002   0x3E94        569               0/0/0

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0001.0000.0001.00-00  0x00000008   0xB20F        573               0/0/0
0001.0000.0001.01-00  0x00000002   0x32A4        317               0/0/0
0001.0000.0002.00-00  0x00000004   0x5A80        568               0/0/0
0001.0000.0003.00-00* 0x0000000A   0xB2CE        979               0/0/0
    0001.0000.0003.01-00* 0x00000002   0x3E94        569               0/0/0

R1#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       --
0001.0000.0002       10         0001.0000.0002       eth1
5254.002a.230a
0001.0000.0003       20         0001.0000.0002       eth1
5254.002a.230a

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       --
0001.0000.0002       10         0001.0000.0002       eth1
5254.002a.230a
0001.0000.0003       20         0001.0000.0002       eth1
5254.002a.230a

R2#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       10         0001.0000.0001       eth1
5254.00dc.0b76
0001.0000.0002       --
0001.0000.0003       10         0001.0000.0003       eth2
5254.00a8.940d

IS-IS paths to level-2 routers
```

```
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       10         0001.0000.0001       eth1
5254.00dc.0b76
0001.0000.0002       --
0001.0000.0003       10         0001.0000.0003       eth2
5254.00a8.940d

R3#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       20         0001.0000.0002       eth2
5254.007e.5ade
0001.0000.0002       10         0001.0000.0002       eth2
5254.007e.5ade
0001.0000.0003       --

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0001.0000.0001       20         0001.0000.0002       eth2
5254.007e.5ade
0001.0000.0002       10         0001.0000.0002       eth2
5254.007e.5ade
0001.0000.0003       --
```

Overload Bit

This section provides examples of configuring the IS-IS overload bit. For detailed information about the commands used in these examples, refer to the *Intermediate System to Intermediate System Command Reference*.

Overview

The expanded use of the overload bit in IS-IS proves beneficial for Internet Service Providers (ISPs) operating both BGP and IS-IS, helping them avoid specific routing anomalies. With the overload bit feature, a router can automatically disable the overload bit when BGP achieves convergence.

IS-IS will deactivate the overload bit if BGP fails to converge within ten minutes and the `wait-for-bgp` option is set in the configuration.

On-Startup Wait-for-BGP

When configure the `set overload-bit wait-for-bgp` after a restart (e.g., using commands like `set-overload-bit on-startup wait-for-bgp` or `set-overload-bit on-startup <5-86400>`), the IS-IS router broadcasts the overload bit status to all its neighbors. This informs them not to route traffic through this router, except for traffic destined for networks directly connected to it.

If an interface is configured with the `no set-overload-bit` command, the router will not advertise or set the overload bit during startup.

Topology

Figure 1-12 represents a sample topology with four routers (R1, R2, R3, and R4) interconnected in a linear sequence.

In this network configuration scenario, these routers collaborate to create a dynamic and resilient routing environment. Through the implementation of both BGP and IS-IS routing protocols, these routers establish connections, exchange

routing information, and optimize network performance. Each router plays a unique role in this network, with specific configurations tailored to ensure smooth communication and fault tolerance.

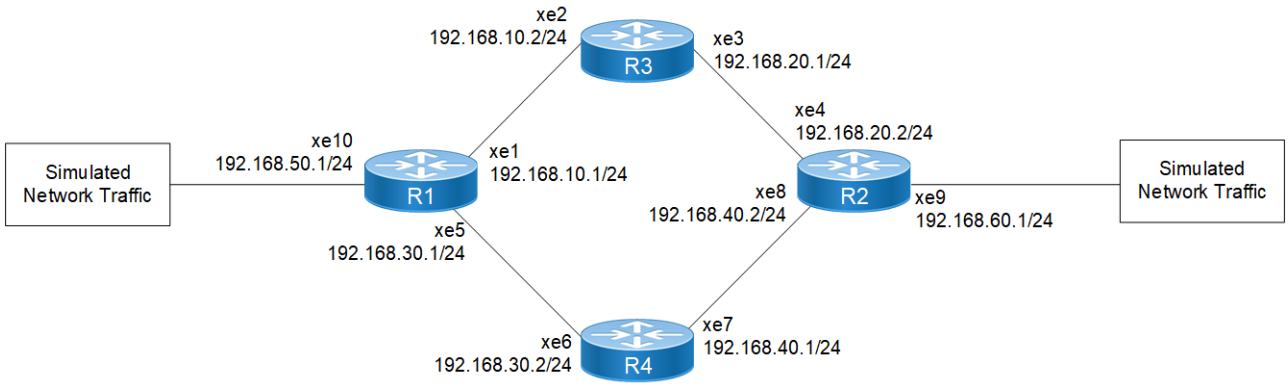

Figure 1-12: Configure Overload Bit Wait-for-BGP

The key element in these configurations is on R3 and R4, where the `set-overload-bit on-startup wait-for-bgp` commands are used. By configuring R3 and R4 with these commands, IS-IS on R3 and R4 will advertise the overload bit to its neighbors, effectively informing them not to route traffic through R3 and R4 until BGP has converged. This behavior helps prevent routing anomalies during network restarts and enhances network stability.

R1

Step 1: Assign IP addresses to the interfaces.

`R1#configure terminal`	Enters configure mode.
`R1(config)#interface xe1`	Configures interface `xe1` and enters interface mode.
`R1(config-if)#ip address 192.168.10.1/24`	Assigns the IP address `192.168.10.1/24`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R1(config)#interface xe5`	Configures interface `xe5` and enters interface mode.
`R1(config-if)#ip address 192.168.30.1/24`	Assigns the IP address `192.168.30.1/24`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R1(config)#interface xe10`	Configures interface `xe10` and enters interface mode.
`R1(config-if)#ip address 192.168.50.1/24`	Assigns the IP address `192.168.50.1/24`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and return to the configure mode.

Step 2: Configure IS-IS operations.

`R1(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R1(config-router)#is-type level-2-only`	Specifies IS-IS to operate as Level-2.
`R1(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.

`R1(config-router)#net 49.0000.0000.0001.00`	Configures IS-IS network entity title.
`R1(config-router)#redistribute connected`	Redistributes connected routes into IS-IS.
`R1(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Configure interfaces to participate in IS-IS process.

`R1(config)#interface xe1`	Configures interface `xe1` and enters interface mode.
`R1(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe1`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R1(config)#interface xe5`	Configures interface `xe5` and enters interface mode.
`R1(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe5`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure a route map and set the weight within the route map.

`R1(config)#route-map rmap1 permit 10`	Configures a route map named `rmap1` with a permit statement and enters route-map mode.
`R1(config-route-map)#set weight 4000`	Sets the weight to 4000 in the route map.
`R1(config-route-map)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-route-map)#exit`	Exits route-map mode and returns to the configure mode.

Step 5: Configure and activate BGP neighbors, and apply the route map to BGP neighbors for both inbound and outbound traffic.

`R1(config)#router bgp 100`	Enters BGP router mode with AS number `100`.
`R1(config-router)#neighbor 192.168.10.2 remote-as 300`	Configures BGP neighbors with remote AS number `300`.
`R1(config-router)#neighbor 192.168.30.2 remote-as 400`	Configures BGP neighbors with remote AS number `400`.
`R1(config-router)#neighbor 192.168.50.2 remote-as 500`	Configures BGP neighbors with remote AS number `500`.
`R1(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R1(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R1(config-router-af)#neighbor 192.168.10.2 activate`	Activates BGP neighbors.
`R1(config-router-af)#neighbor 192.168.10.2 route-map rmap1 in`	Applies the route-map `rmap1` to BGP neighbors for inbound traffic.
`R1(config-router-af)#neighbor 192.168.10.2 route-map rmap1 out`	Applies the route-map `rmap1` to BGP neighbors for outbound traffic.

R1(config-router-af)#neighbor 192.168.30.2 activate	Activates BGP neighbors.
R1(config-router-af)#neighbor 192.168.50.2 activate	Activates BGP neighbors.
R1(config-router-af)#exit-address-family	Exits BGP address family mode.
R1(config-router)#commit	Commits the candidate configuration to the running configuration.
R1(config-router)#exit	Exits BGP router mode and returns to the configure mode.

R2

Step 1: Assign IP addresses to the interfaces.

R2#configure terminal	Enters configure mode.
R2(config)#interface xe4	Configures interface xe4 and enters interface mode.
R2(config-if)#ip address 192.168.20.2/24	Assigns the IP address 192.168.20.2/24.
R2(config-if)#commit	Commits the candidate configuration to the running configuration.
R2(config-if)#exit	Exits interface mode and returns to the configure mode.
R2(config)#interface xe8	Configures interface xe8 and enters interface mode.
R2(config-if)#ip address 192.168.40.2/24	Assigns the IP address 192.168.40.2/24.
R2(config-if)#commit	Commits the candidate configuration to the running configuration.
R2(config-if)#exit	Exits interface mode and returns to the configure mode.
R2(config)#interface xe9	Configures interface xe9 and enters interface mode.
R2(config-if)#ip address 192.168.60.1/24	Assigns the IP address 192.168.60.1/24.
R2(config-if)#commit	Commits the candidate configuration to the running configuration.
R2(config-if)#exit	Exits interface mode and return to the configure mode.

Step 2: Configure IS-IS IPv4 routing.

R2(config)#router isis test	Enters IS-IS router mode for IS-IS process named test.
R2(config-router)#is-type level-2-only	Specifies IS-IS to operate as Level-2.
R2(config-router)#bfd all-interfaces	Enables BFD on all interfaces.
R2(config-router)#net 50.0000.0000.0002.00	Configures IS-IS network entity title.
R2(config-router)#redistribute connected	Redistributes connected routes into IS-IS.
R2(config-router)#commit	Commits the candidate configuration to the running configuration.
R2(config-router)#exit	Exits IS-IS router mode and returns to the configure mode.

Step 3: Enable IS-IS IPv4 routing on the interfaces.

R2(config)#interface xe4	Configures interface xe4 and enters interface mode.
R2(config-if)#ip router isis test	Enables IS-IS IPv4 routing on the interface xe4.

`R2(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R2(config)#interface xe8`	Configures interface `xe8` and enters interface mode.
`R2(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe8`.
`R2(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure a route map and set the weight within the route map.

`R2(config)#route-map rmap1 permit 10`	Configures a route map named `rmap1` with a permit statement and enters route-map mode.
`R2(config-route-map)#set weight 4000`	Sets the weight to `4000` in the route map.
`R2(config-route-map)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-route-map)#exit`	Exits route-map mode and returns to the configure mode.

Step 5: Configure and activate BGP neighbors, and apply the route map to BGP neighbors for both inbound and outbound traffic.

`R2(config)#router bgp 200`	Enters BGP router mode with AS number `200`.
`R2(config-router)#neighbor 192.168.20.1 remote-as 300`	Configures BGP neighbors with remote AS number `300`.
`R2(config-router)#neighbor 192.168.40.1 remote-as 400`	Configures BGP neighbors with remote AS number `400`.
`R2(config-router)#neighbor 192.168.60.2 remote-as 600`	Configures BGP neighbors with remote AS number `600`.
`R2(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R2(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R2(config-router-af)#neighbor 192.168.20.1 activate`	Activates BGP neighbors.
`R2(config-router-af)#neighbor 192.168.20.1 route-map rmap1 in`	Applies the route-map `rmap1` to BGP neighbors for inbound traffic.
`R2(config-router-af)#neighbor 192.168.20.1 route-map rmap1 out`	Applies the route-map `rmap1` to BGP neighbors for outbound traffic.
`R2(config-router-af)#neighbor 192.168.40.1 activate`	Activates BGP neighbors.
`R2(config-router-af)#neighbor 192.168.60.2 activate`	Activates BGP neighbors.
`R2(config-router-af)#exit-address-family`	Exits BGP address family mode.
`R2(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-router)#exit`	Exits BGP router mode and returns to the configure mode.

R3

Step 1: Assign IP addresses to the interfaces.

`R3#configure terminal`	Enters configure mode.
`R3(config)#interface xe2`	Configures interface `xe2` and enters interface mode.
`R3(config-if)#ip address 192.168.10.2/24`	Assigns the IP address `192.168.10.2/24`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R3(config)#interface xe3`	Configures interface `xe3` and enters interface mode.
`R3(config-if)#ip address 192.168.20.1/24`	Assigns the IP address `192.168.20.1/24`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS IPv4 routing, setting the overload bit on startup and waiting for BGP convergence.

`R3(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R3(config-router)#is-type level-2-only`	Specifies IS-IS to operate as Level-2.
`R3(config-router)#set-overload-bit on-startup wait-for-bgp`	Configures the overload bit to be set on startup and wait for BGP convergence.
`R3(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.
`R3(config-router)#net 51.0000.0000.0003.00`	Configures IS-IS network entity title.
`R3(config-router)#redistribute connected`	Redistributes connected routes into IS-IS.
`R3(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Enable IS-IS IPv4 routing on the interfaces.

`R3(config)#interface xe2`	Configures interface `xe2` and enters interface mode.
`R3(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe2`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R3(config)#interface xe3`	Configures interface `xe3` and enters interface mode.
`R3(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe3`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure and activate BGP neighbors.

`R3(config)#router bgp 300`	Enters BGP router mode with AS number `300`.
`R3(config-router)#neighbor 192.168.10.1 remote-as 100`	Configures BGP neighbors with remote AS number `100`.

`R3(config-router)#neighbor 192.168.20.2 remote-as 200`	Configures BGP neighbors with remote AS number `200`.
`R3(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R3(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R3(config-router-af)#neighbor 192.168.10.1 activate`	Activates BGP neighbors.
`R3(config-router-af)#neighbor 192.168.20.2 activate`	Activates BGP neighbors.
`R3(config-router-af)#exit-address-family`	Exits BGP address family mode.
`R3(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-router)#exit`	Exits BGP router mode and returns to the configure mode.

R4

Step 1: Assign IP addresses to the interfaces.

`R4#configure terminal`	Enters configure mode.
`R4(config)#interface xe6`	Configures interface `xe6` and enters interface mode.
`R4(config-if)#ip address 192.168.30.2/24`	Assigns the IP address `192.168.30.2/24`.
`R4(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R4(config)#interface xe7`	Configures interface `xe7` and enters interface mode.
`R4(config-if)#ip address 192.168.40.1/24`	Assigns the IP address `192.168.40.1/24`.
`R4(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS IPv4 routing, setting the overload bit on startup and waiting for BGP convergence.

`R4(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R4(config-router)#is-type level-2-only`	Specifies IS-IS to operate as Level-2.
`R4(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.
`R4(config-router)#net 52.0000.0000.0004.00`	Configures IS-IS network entity title.
`R4(config-router)#redistribute connected`	Redistributes connected routes into IS-IS.
`R4(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Enable IS-IS IPv4 routing on the interfaces.

`R4(config)#interface xe6`	Configures interface `xe6` and enters interface mode.
`R4(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe6`.

`R4(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R4(config)#interface xe7`	Configures interface `xe7` and enters interface mode.
`R4(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe7`.
`R4(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure and activate BGP neighbors.

`R4(config)#router bgp 400`	Enters BGP router mode with AS number `400`.
`R4(config-router)#neighbor 192.168.30.1 remote-as 100`	Configures BGP neighbors with remote AS number `100`.
`R4(config-router)#neighbor 192.168.40.2 remote-as 200`	Configures BGP neighbors with remote AS number `200`.
`R4(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R4(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R4(config-router-af)#neighbor 192.168.30.1 activate`	Activates BGP neighbors.
`R4(config-router-af)#neighbor 192.168.40.2 activate`	Activates BGP neighbors.
`R4(config-router-af)#exit-address-family`	Exits BGP address family mode.
`R4(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R4(config-router)#exit`	Exits BGP router mode and returns to the configure mode.

Validation

1. Verify the ISIS and BGP sessions are coming up.
2. Send traffic from one network simulator to another and verify traffic flow via R1-R3-R2 router's without loss.
3. Reload R3 router.
4. Verify traffic switchover to the R1-R4-R2 routers path.
5. Unconfigure the BGP neighbor between R2 and R3 routers during R2 routers reload.
6. After reloading, check the overload-bit is set to `1 (0/0/1)` on the R3 router using the `show isis database` command.
7. Now configure the BGP neighbor between R2 and R3 routers and verify the BGP session comes up.
8. After some time, verify the overload-bit is set to `0 (0/0/0)` using the `show isis database` command on the R3 router.
9. Verify traffic switchover to the R1-R3-R2 router's path.

Before Reload

```
R3#show interface counters rate mbps
```

```
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 xe4                 680.47        664526       0.01          4
 xe5                 0.00          4            680.48        664525

R3#reload
Are you sure you would like to reset the system? (y/n): y
 Device rebooted by ocnos user from cmlsh at time:Thu Dec  7 11:45:40 2023

R3#         Stopping xinetd service...
[  OK  ] Stopped target Timers.
[  OK  ] Stopped Daily rotation of log files.
[  OK  ] Stopped Daily Cleanup of Temporary Directories.
```

After Reload

R3

```
R3#show interface counters rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 xe4                 0.00          4            0.01          5
 xe5                 0.00          4            0.01          4

R3#show isis  database
Tag test:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000057   0xDE86        1130              0/0/0
0000.0000.0002.00-00  0x00000056   0x7E6E        1130              0/0/0
0000.0000.0003.00-00* 0x00000055   0x55C5        1131              0/0/1
0000.0000.0003.01-00* 0x00000005   0x25AD        1132              0/0/0
0000.0000.0003.02-00* 0x0000004A   0x7A13        1132              0/0/0
0000.0000.0004.00-00  0x00000051   0x20C2        574               0/0/0
0000.0000.0004.01-00  0x0000004C   0x9CEC        466               0/0/0
0000.0000.0004.02-00  0x00000049   0x820A        467               0/0/0
```

R4

```
R4#show interface counters rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 ge1                 0.00          4            691.93        675712
 xe8                 691.93        675712       0.01          4
```

R2

```
R2(config-router)#no  neighbor 192.168.20.1 remote-as 300
R2(config-router)#commit
R2#
R2#config t
R2(config)#router bgp 200
R2(config-router)#neighbor 192.168.20.1 remote-as 300
```

```
R2(config-router)#address-family ipv4 unicast
R2(config-router-af)#neighbor 192.168.20.1 activate
R2(config-router-af)#neighbor 192.168.20.1 route-map rmap1 in
R2(config-router-af)#neighbor 192.168.20.1 route-map rmap1 out
R2(config-router-af)#exit-address-family
```

R3

```
R3#show isis  database
Tag test:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000057   0xDE86        1092              0/0/0
0000.0000.0002.00-00  0x00000056   0x7E6E        1092              0/0/0
0000.0000.0003.00-00* 0x00000056   0x4FCE        1191              0/0/0
0000.0000.0003.01-00* 0x00000005   0x25AD        1094              0/0/0
0000.0000.0003.02-00* 0x0000004A   0x7A13        1094              0/0/0
0000.0000.0004.00-00  0x00000051   0x20C2        536               0/0/0
0000.0000.0004.01-00  0x0000004C   0x9CEC        428               0/0/0
0000.0000.0004.02-00  0x00000049   0x820A        429               0/0/0

R3#show interface counters rate mbps
+-------------------+-------------+-------------+-------------+-------------+
|     Interface     |    Rx mbps  |    Rx pps   |    Tx mbps  |    Tx pps   |
+-------------------+-------------+-------------+-------------+-------------+
 xe4                  681.59        665618        0.01          4
 xe5                  0.00          4             681.60        665618
```

Suppress External and Interlevel Reachability

The control mechanism of suppressing the redistribution of external, interlevel, or both external and interlevel reachabilities is useful for managing the advertisement of these specific reachability data during overload states, providing flexibility in network configurations.

Topology

Figure 1-13 represents a sample topology with three routers (R1, R2, and R3) interconnected in a linear sequence.

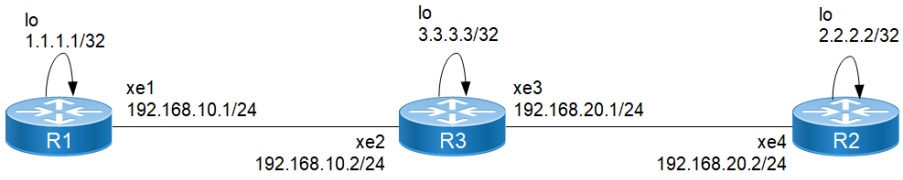

Figure 1-13: Configure Overload Bit Suppress

The key element in the below configurations is on R3 router, where the `set-overload-bit suppress external` and `set-overload-bit suppress interlevel` commands are used.

Configuring the R3 router with the `set-overload-bit suppress external` or `set-overload-bit suppress interlevel` commands in IS-IS prevents the advertisement of external routes learned from BGP to its ISIS neighbors or prevents the advertisement of interlevel routes to its ISIS neighbors. The options `suppress external` or `suppress interlevel` contribute to the controlled management of routing information during network overload states.

R1

Step 1: Assign IP addresses to the interfaces.

`R1#configure terminal`	Enters configure mode.
`R1(config)#interface xe1`	Configures interface `xe1` and enters interface mode.
`R1(config-if)#ip address 192.168.10.1/24`	Assigns the IP address `192.168.10.1/24`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS operations.

`R1(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R1(config-router)#is-type level-1`	Specifies IS-IS to operate as Level-1.
`R1(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.
`R1(config-router)#net 49.0000.0000.0001.00`	Configures IS-IS network entity title.
`R1(config-router)#redistribute connected`	Redistributes connected routes into IS-IS.
`R1(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Configure interfaces to participate in IS-IS process.

`R1(config)#interface xe1`	Configures interface `xe1` and enters interface mode.
`R1(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe1`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R1(config)#interface lo`	Configures a loopback interface `lo` and enters interface mode.
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Assigns the secondary IP address `1.1.1.1/32`.
`R1(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the loopback interface `lo`.
`R1(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure and activate BGP neighbors

Note: Configure BGP only for Suppress External reachability using the `set-overload-bit suppress external` command, excluding interlevel reachability.

`R1(config)#router bgp 100`	Enters BGP router mode with AS number `100`.
`R1(config-router)#neighbor 192.168.10.2 remote-as 300`	Configures BGP neighbors with remote AS number `300`.

`R1(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R1(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R1(config-router-af)#neighbor 192.168.10.2 activate`	Activates BGP neighbors.
`R1(config-router-af)#exit-address-family`	Exits BGP address family mode.
`R1(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exits BGP router mode and returns to the configure mode.

R2

Step 1: Assign IP addresses to the interfaces.

`R2#configure terminal`	Enters configure mode.
`R2(config)#interface xe4`	Configures interface `xe4` and enters interface mode.
`R2(config-if)#ip address 192.168.20.2/24`	Assigns the IP address `192.168.20.2/24`.
`R2(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS IPv4 routing.

`R2(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R2(config-router)#is-type level-1`	Specifies IS-IS to operate as Level-1.
`R2(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.
`R2(config-router)#net 49.0000.0000.0002.00`	Configures IS-IS network entity title.
`R2(config-router)#redistribute connected`	Redistributes connected routes into IS-IS.
`R2(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Enable IS-IS IPv4 routing on the interfaces.

`R2(config)#interface xe4`	Configures interface `xe4` and enters interface mode.
`R2(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe4`.
`R2(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R2(config)#interface lo`	Configures a loopback interface `lo` and enters interface mode.
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Assigns the secondary IP address `2.2.2.2/32`.
`R2(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the loopback interface `lo`.

`R2(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R2(config-if)#exit`	Exits interface mode and returns to the configure mode.

R3

Step 1: Assign IP addresses to the interfaces

`R3#configure terminal`	Enters configure mode.
`R3(config)#interface xe2`	Configures interface `xe2` and enters interface mode.
`R3(config-if)#ip address 192.168.10.2/24`	Assigns the IP address `192.168.10.2/24`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R3(config)#interface xe3`	Configures interface `xe3` and enters interface mode.
`R3(config-if)#ip address 192.168.20.1/24`	Assigns the IP address `192.168.20.1/24`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS IPv4 routing, setting the overload bit suppress external or interlevel options.

`R3(config)#router isis test`	Enters IS-IS router mode for IS-IS process named `test`.
`R3(config-router)#is-type level-2-only`	Specifies IS-IS to operate as Level-2.
`R3(config-router)#set-overload-bit suppress external` OR **`R3(config-router)#set-overload-bit suppress interlevel`**	Suppresses the redistribution of external reachability during overload states. OR Suppresses the redistribution of interlevel reachability during overload states.
`R3(config-router)#bfd all-interfaces`	Enables BFD on all interfaces.
`R3(config-router)#net 49.0000.0000.0003.00`	Configures the IS-IS network entity title.
`R3(config-router)#redistribute bgp`	Redistributes BGP routes into IS-IS.
`R3(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-router)#exit`	Exits IS-IS router mode and returns to the configure mode.

Step 3: Enable IS-IS IPv4 routing on the interfaces

`R3(config)#interface xe2`	Configures interface `xe2` and enters interface mode.
`R3(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe2`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R3(config)#interface xe3`	Configures interface `xe3` and enters interface mode.
`R3(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the interface `xe3`.

`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.
`R3(config)#interface lo`	Configures a loopback interface `lo` and enters interface mode.
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Assigns the secondary IP address `3.3.3.3/32`.
`R3(config-if)#ip router isis test`	Enables IS-IS IPv4 routing on the loopback interface `lo`.
`R3(config-if)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 4: Configure and activate BGP neighbors

Note: Configure BGP only for Suppress External reachability using the `set-overload-bit suppress external` command, excluding interlevel reachability.

`R3(config)#router bgp 300`	Enters BGP router mode with AS number `300`.
`R3(config-router)#neighbor 192.168.10.1 remote-as 100`	Configures BGP neighbors with remote AS number `100`.
`R3(config-router)#address-family ipv4 unicast`	Enters BGP address family mode for IPv4 unicast.
`R3(config-router-af)#redistribute connected`	Redistributes connected routes into BGP.
`R3(config-router-af)#neighbor 192.168.10.1 activate`	Activates BGP neighbors.
`R3(config-router-af)#exit-address-family`	Exits BGP address family mode.
`R3(config-router)#commit`	Commits the candidate configuration to the running configuration.
`R3(config-router)#exit`	Exits BGP router mode and returns to the configure mode.

Validation

Verify all IS-IS sessions are up.

Pre-Configuration Check: ISIS Session Status

Before configuring `set-overload-bit suppress external` command, verify the ISIS session status on R1, R2, and R3 routers.

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:13:54
i L1         2.2.2.2/32 [115/20] via 192.168.10.2, xe2, 00:01:46
```

```
i L1          3.3.3.3/32 [115/30] via 192.168.10.2, xe2, 00:00:06
C             127.0.0.0/8 is directly connected, lo, 00:13:54
C             192.168.10.0/24 is directly connected, xe1, 00:13:54
i L1          192.168.20.0/24 [115/20] via 192.168.10.2, xe2, 00:01:46

Gateway of last resort is not set
```

```
R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1          1.1.1.1/32 [115/30] via 192.168.10.1, xe1, 00:01:19
i L1          2.2.2.2/32 [115/20] via 192.168.20.2, xe4, 00:06:49
C             3.3.3.3/32 is directly connected, lo, 00:25:23
C             127.0.0.0/8 is directly connected, lo, 00:25:23
i L1          192.168.10.0/24 is directly connected, xe2, 00:06:49
C             192.168.20.0/24 is directly connected, xe3, 00:25:23

Gateway of last resort is not set
```

```
R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B             1.1.1.1/32 [20/0] via 192.168.20.1, xe3, 00:00:05
C             2.2.2.2/32 is directly connected, lo, 00:25:20
i L1          3.3.3.3/32 [115/20] via 192.168.20.1, xe3, 00:06:46
C             127.0.0.0/8 is directly connected, lo, 00:25:20
C             192.168.10.0/24 [100/23] via 192.168.20.1, xe3, 00:25:20
C             192.168.20.0/24 is directly connected, xe4, 00:25:20

Gateway of last resort is not set
```

Before configuring `set-overload-bit suppress interlevel` command, verify the ISIS session status on R1, R2, and R3 routers.

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:14:58
i L1         2.2.2.2/32 [115/20] via 192.168.10.2, xe2, 00:02:50
i L1         3.3.3.3/32 [115/30] via 192.168.10.2, xe2, 00:01:10
C            127.0.0.0/8 is directly connected, lo, 00:14:58
C            192.168.10.0/24 is directly connected, xe1, 00:14:58
i L1         192.168.20.0/24 [115/20] via 192.168.10.2, xe2, 00:02:50

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/30] via 192.168.10.1, xe1, 00:01:19
i L1         2.2.2.2/32 [115/20] via 192.168.20.21, xe4, 00:06:49
C            3.3.3.3/32 is directly connected, lo, 00:25:23
C            127.0.0.0/8 is directly connected, lo, 00:25:23
i L1         192.168.10.0/24 is directly connected, xe2, 00:06:49
C            192.168.20.0/24 is directly connected, xe3, 00:25:23

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/20] via 192.168.20.1, xe3, 00:06:41
```

```
C            2.2.2.2/32 is directly connected, lo, 00:25:15
i L1         3.3.3.3/32 [115/20] via 192.168.20.1, xe3, 00:06:41
C            127.0.0.0/8 is directly connected, lo, 00:25:15
C            192.168.10.0/24 [115/20] via 192.168.20.1, xe3, 00:25:15
C            192.168.20.0/24 is directly connected, xe4, 00:25:15

Gateway of last resort is not set
```

Post-Configuration Check: ISIS Session Status

After configuring `set-overload-bit suppress external` command, verify the ISIS session status on R1, R2, and R3 routers.

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:20:29
i L1         2.2.2.2/32 [115/20] via 192.168.10.2, xe2, 00:01:56
C            127.0.0.0/8 is directly connected, lo, 00:20:29
C            192.168.10.0/24 is directly connected, xe1, 00:20:29
i L1         192.168.20.0/24 [115/20] via 192.168.10.2, xe2, 00:01:56

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         2.2.2.2/32 [115/20] via 192.168.20.2, xe4, 00:07:50
C            3.3.3.3/32 is directly connected, lo, 00:26:24
C            127.0.0.0/8 is directly connected, lo, 00:26:24
i L1         192.168.10.0/24 is directly connected, xe2, 00:07:50
C            192.168.20.0/24 is directly connected, xe3, 00:26:24

Gateway of last resort is not set

R2#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
B            1.1.1.1/32 [20/0] via 192.168.20.1, xe3, 00:00:03
C            2.2.2.2/32 is directly connected, lo, 00:25:18
i L1         3.3.3.3/32 [115/20] via 192.168.20.1, xe3, 00:06:44
C            127.0.0.0/8 is directly connected, lo, 00:25:18
C            192.168.10.0/24 [100/10] via 192.168.20.1, xe3, 00:25:18
C            192.168.20.0/24 is directly connected, xe4, 00:25:18

Gateway of last resort is not set
```

After configuring `set-overload-bit suppress interlevel` command, verify the ISIS session status on R1, R2, and R3 routers.

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:14:58
i L1         2.2.2.2/32 [115/20] via 192.168.10.2, xe2, 00:02:50
C            127.0.0.0/8 is directly connected, lo, 00:14:58
C            192.168.10.0/24 is directly connected, xe1, 00:14:58
i L1         192.168.20.0/24 [115/20] via 192.168.10.2, xe2, 00:02:50

Gateway of last resort is not set
```

```
R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
```

```
i L1          1.1.1.1/32 [115/30] via 192.168.10.1, xe1, 00:01:19
i L1          2.2.2.2/32 [115/20] via 192.168.20.2, xe4, 00:06:49
C             3.3.3.3/32 is directly connected, lo, 00:25:23
C             127.0.0.0/8 is directly connected, lo, 00:25:23
i L1          192.168.10.0/24 is directly connected, xe2, 00:06:49
C             192.168.20.0/24 is directly connected, xe3, 00:25:23

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C             2.2.2.2/32 is directly connected, lo, 00:25:15
i L1          3.3.3.3/32 [115/20] via 192.168.20.1, xe3, 00:06:41
C             127.0.0.0/8 is directly connected, lo, 00:25:15
C             192.168.10.0/24 [115/20] via 192.168.20.1, xe3, 00:25:15
C             192.168.20.0/24 is directly connected, xe4, 00:25:15

    Gateway of last resort is not set
```

Passive Interface

In ISP and large enterprise networks, many of the distribution routers have more than 200 interfaces. Before the Default Passive-Interface feature, there were two possibilities for obtaining routing information from all of these interfaces:

- Configure a routing protocol on the backbone interfaces and redistribute connected interfaces.
- Configure the routing protocol on all interfaces and manually set most of them as passive, which was time consuming.

The solution to this problem was to configure the routing protocol on all interfaces and manually set the `passive-interface` command on the interfaces where adjacency was not desired. In certain networks, this meant coding 200 or more passive-interface statements. With the Default Passive Interface feature, this problem is solved by allowing all interfaces to be set as passive by default using a single `passive-interface` default command, then configuring individual interfaces in which adjacencies are desired using the `passive-interface <interface-name> disable` command.

Usage

1. When a specific interface is configured as passive using the `passive-interface <interface-name>` command:
    - The interface loses its adjacency on that interface, for example, eth1.

- The interface (eth1) is still advertised by other IS-IS speaking interfaces to their neighbors.

2. When a specific interface is configured as passive using `passive-interface <interface-name>` command followed by removing the configuration using `no passive-interface <interface-name>` command:
   - The interface is IS-IS disabled and must be enabled using the `ip router isis` command (for example, `ip router isis 1`).
   - If IS-IS is not configured, the interface (for example, eth1) will not be advertised by other IS-IS speaking interfaces to their neighbors.
3. When an interface is configured with the `passive interface` command:
   - All IS-IS enabled interfaces lose their adjacency.
   - All IS-IS enabled interfaces in the system will be made passive.
   - To establish adjacency on a particular interface, the `passive interface <interface-name> disable` command must be enabled.
   - All interfaces which were made passive are advertised by the active IS-IS speaking interface to its neighbors.
4. When an interface is configured with the `no passive interface` command:
   - All interfaces which are currently passive, will become active.
   - If IS-IS is configured on those interface, it will start sending out IS-IS packets and attempt to form adjacency.
   - If IS-IS is not configured on those interfaces, it will not be advertised by the active IS-IS speaking interface to its neighbors.

Topology

Figure 1-14 shows a passive-interface configuration example.

Figure 1-14: IS-IS Passive Interface

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance `(1)`.
`(config-router)#net 49.0001.0000.0000.0001.00`	Define the NET address.
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure and enter Interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface `eth1` (connected to R2).

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Specify the interface (eth1)to configure and enter Interface mode.
`(config-if)#ip address 20.20.20.2/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R1).
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#interface eth2`	Specify the interface (eth2)to configure and enter Interface mode.
`(config-if)#ip address 30.30.30.1/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth2 (connected to R3).
`(config-if)#exit`	Exit interface mode and return to Configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#net 49.0001.0000.0000.0002.00`	Define the NET address.
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#passive-interface eth1`	Configure the eth1 interface as passive.

R3

`#configure terminal`	Enter configure mode.
`(config)#router isis 1`	Create an IS-IS routing instance (1).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0001.0000.0000.0003.00`	Define the NET address.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth2`	Specify the interface (eth2)to configure and enter Interface mode.
`(config-if)#ip address 30.30.30.2/24`	Configure IP address on interface.
`(config-if)#ip router isis 1`	Enable IS-IS routing on interface eth1 (connected to R2).

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 0
Total number of adjacencies: 0
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 1:  VRF : default
```

```
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0003 eth2        5254.00a8.940d  Up     9         L2   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.007e.5ade  Up     19        L2   IS-IS

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       10          --                  eth1            0

R2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    20.20.20.0/24       0           --                  eth1            0
C    30.30.30.0/24       10          --                  eth2            0

R3#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L2   20.20.20.0/24       10          30.30.30.1          eth2            0
C    30.30.30.0/24       10          --                  eth2            0

R1#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000004   0x3A02        923               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.1
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0001.01-00* 0x00000001   0xF108        0 (923)           0/0/0

R2#show isis database verbose
Tag 1:  VRF : default
```

```
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000002   0x3EFF        887               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.1
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0001.01-00  0x00000001   0x21B9        888               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000003   0x3761        906               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   0          IP 20.20.20.0 255.255.255.0
  Metric:   10         IP 30.30.30.0 255.255.255.0
0000.0000.0003.00-00  0x00000002   0x530E        909               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IP 30.30.30.0 255.255.255.0
0000.0000.0003.01-00  0x00000001   0x2DA9        905               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00

R3#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000002   0x3EFF        883               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   20.20.20.1
  Metric:   10         IP 20.20.20.0 255.255.255.0
0000.0000.0001.01-00  0x00000001   0x21B9        884               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00  0x00000003   0x3761        901               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   0          IP 20.20.20.0 255.255.255.0
  Metric:   10         IP 30.30.30.0 255.255.255.0
0000.0000.0003.00-00* 0x00000002   0x530E        906               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   30.30.30.2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IP 30.30.30.0 255.255.255.0
0000.0000.0003.01-00* 0x00000001   0x2DA9        902               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00
```

```
R1#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --

R2#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       **
0000.0000.0002       --
0000.0000.0003       10         0000.0000.0003       eth2
5254.00a8.940d

R3#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       **
0000.0000.0002       10         0000.0000.0002       eth2
5254.007e.5ade
0000.0000.0003       --
```

IS-IS IPv4 Loop-Free Alternate Fast Reroute

This section contains IS-IS (Intermediate System to Intermediate System) Loop-Free Alternate Fast Reroute (LFA-FRR) configuration examples.

For details about the commands used in these examples, see the *Intermediate System to Intermediate System Command Reference*.

Overview

When a primary next-hop fails, LFA-FRR reduces the failure reaction time to tens of milliseconds using a pre-computed alternate next- hop, so that the alternate can be rapidly used when the failure is detected. A network with this feature experiences less traffic loss and less micro-looping of packets than a network without LFA-FRR.

After you enable LFA-FRR, routers calculate a backup path for each primary path to reach the destination. The backup path is calculated based on the attributes such as node protecting, link protecting, and broadcast link protecting. If there is an ECMP path to reach prefixes, the backup is selected from the same primary set by default; if a secondary tie-breaker is enabled, and if a secondary path is available, the backup will be selected from the secondary path.

Basic Configuration

Topology

Figure 1-15 shows the configuration to enable the basic ISIS LFA feature.

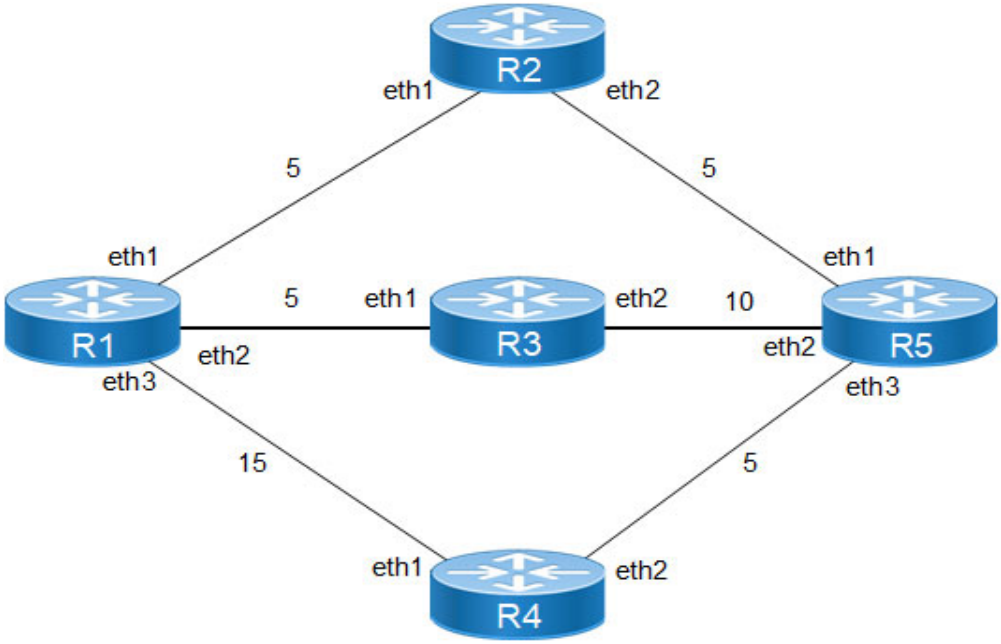

Figure 1-15: ISIS LFA-FRR

R1

`#configure terminal`	Enter configure mode.
`(config)#int eth1`	Enter interface mode.
`(config-if)#ip address 10.10.10.142/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#isis metric 5`	Configure isis metric value for interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth2`	Enter interface mode.
`(config-if)#ip address 20.20.20.142/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#isis metric 5`	Configure isis metric value for interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth3`	Enter interface mode.
`(config-if)#ip address 30.30.30.142/24`	Configure the IP address of the interface.

(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#isis metric 15	Configure isis metric value for interface
(config-if)#exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 with instance 1
(config-router)#net 49.0000.0000.0001.00	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for ISIS on all interfaces
(config-router)#fast-reroute per-prefix level-1 proto ipv4 all	Configure LFA-FRR to calculate the available backup path for all L1 ipv4 prefixes learnt
(config-router)#exit	Exit router mode.
(config)#exit	Exit config mode.

R2

#configure terminal	Enter configure mode.
(config)#int eth1	Enter interface mode.
(config-if)#ip address 10.10.10.141/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#isis metric 5	Configure isis metric value for interface
(config-if)#exit	Exit interface mode.
(config)#int eth2	Enter interface mode.
(config-if)#ip address 40.40.40.141/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#isis metric 5	Configure isis metric value for interface
(config-if)#exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 with instance 1
(config-router)#net 49.0000.0000.0002.00	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for ISIS on all interfaces
(config-router)#exit	Exit router mode.
(config)#exit	Exit config mode.

R3

#configure terminal	Enter configure mode.
(config)#int eth1	Enter interface mode.
(config-if)#ip address 20.20.20.143/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1

(config-if)#isis metric 5	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#int eth2	Enter interface mode.
(config-if)#ip address 50.50.50.143/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 10	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 with instance 1
(config-router)#net 49.0000.0000.0003.00	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for ISIS on all interfaces
(config-router)#exit	Exit router mode.
(config)#exit	Exit config mode.

R4

#configure terminal	Enter configure mode.
(config)#int eth1	Enter interface mode.
(config-if)#ip address 30.30.30.144/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 15	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#int eth2	Enter interface mode.
(config-if)#ip address 60.60.60.144/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 5	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 with instance 1
(config-router)#net 49.0000.0000.0004.00	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for ISIS on all interfaces
(config-router)#exit	Exit router mode.
(config)#exit	Exit config mode.

R5

`#configure terminal`	Enter configure mode.
`(config)#int eth1`	Enter interface mode.
`(config-if)#ip address 40.40.40.145/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth2`	Enter interface mode.
`(config-if)#ip address 50.50.50.145/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth3`	Enter interface mode.
`(config-if)#ip address 60.60.60.145/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area `49` with instance 1
`(config-router)#net 49.0000.0000.0005.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

R1

ISIS neighborship:

```
R1#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     27        L1   IS-IS
0000.0000.0003 eth2        5254.00dc.2f11  Up     7         L1   IS-IS
0000.0000.0004 eth3        5254.00f5.35a4  Up     7         L1   IS-IS
```

```
R2#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00dc.0b76  Up     6         L1   IS-IS
0000.0000.0005 eth2        5254.00b3.110c  Up     7         L1   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00a1.6afe  Up     22        L1   IS-IS
0000.0000.0005 eth2        5254.0056.7a3d  Up     27        L1   IS-IS

R4#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.0011.a028  Up     21        L1   IS-IS
0000.0000.0005 eth2        5254.00d3.fb41  Up     21        L1   IS-IS


R5#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.007e.5ade  Up     27        L1   IS-IS
0000.0000.0003 eth2        5254.00a8.940d  Up     6         L1   IS-IS
0000.0000.0004 eth3        5254.00e2.aece  Up     7         L1   IS-IS
```

Check the ISIS route installation with primary and backup paths in the ISIS table and RIB table.

Primary paths:

```
R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    10.10.10.0/24       5           --                  eth1            0
C    20.20.20.0/24       5           --                  eth2            0
C    30.30.30.0/24       15          --                  eth3            0
```

```
L1   40.40.40.0/24       10          10.10.10.141        eth1            0
L1   50.50.50.0/24       15          20.20.20.143        eth2            0
L1   60.60.60.0/24       15          10.10.10.141        eth1            0

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth1, 00:43:14
C            10.12.30.0/24 is directly connected, eth0, 01:42:55
C            20.20.20.0/24 is directly connected, eth2, 00:43:14
C            30.30.30.0/24 is directly connected, eth3, 00:43:14
i L1         40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:16:42
i L1         50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:16:55
i L1         60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:16:42
C            127.0.0.0/8 is directly connected, lo, 01:42:55

Gateway of last resort is not set
```

R1#FRR backup paths:

```
R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  40.40.40.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 25
    Protection Provided : LP NP BP

L1  50.50.50.0/24
    Primary Path via    : 20.20.20.143, eth2
    FRR Backup Path via : 10.10.10.141, eth1
    FRR Metric          : 20
    Protection Provided : LP NP BP

L1  60.60.60.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 20
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info
       * - candidate default

IP Route Table for VRF "default"
i L1    40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:18:01
                      [FRR-NH] via 30.30.30.144, eth3

i L1    50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:18:14
                      [FRR-NH] via 10.10.10.141, eth1

i L1    60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:18:01
                      [FRR-NH] via 30.30.30.144, eth3
```

It is not mandatory that for all primary paths, there exists an LFA backup path only if inequality equation satisfies according to attributes configured on routers, backup path will be calculated.

To prohibit an interface from being used as a repair path, disable fast reroute calculation on the interface:

`R1(config-if)#interface eth1`	Enter interface.
`(config-if)#isis fast-reroute per-prefix candidate disable level-1`	Disable fast reroute calculation on the interface.
`(config-if)#end`	Exit.

Verify that the `eth1` interface is not used for backup path calculation.

```
R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  40.40.40.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 25
    Protection Provided : LP NP BP

L1  50.50.50.0/24
    Primary Path via    : 20.20.20.143, eth2
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 30
    Protection Provided : LP NP BP

L1  60.60.60.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 20
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info
       * - candidate default

IP Route Table for VRF "default"
i L1    40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:20:22
                      [FRR-NH] via 30.30.30.144, eth3

i L1    50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:20:35
                      [FRR-NH] via 30.30.30.144, eth3

i L1    60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:20:22
                      [FRR-NH] via 30.30.30.144, eth3

R1#
```

Backup Path based on Route-Map Prefixes

R1

Route-map and access-list configuration:

`(config)#ip access-list 1`	Create an access-list
`(config-ip-acl)#permit any 40.40.40.0 0.0.0.255 any`	Configuring rule to permit only one prefix
`(config)#route-map rmap1 permit 1`	Create a route-map
`(config-route-map)#match ip address 1`	Apply above created access-list in route-map
`(config)#exit`	Exit config mode.

Apply the above created route-map with fast-reroute:

`(config)#router isis 1`	Create an IS-IS routing instance for area `49` with instance 1
`(config-router)#fast-reroute per-prefix level-1 proto ipv4 route-map rmap1`	Configure LFA-FRR to calculate the available backup path for routes allowed through route-map
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
```

```
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth1, 01:02:04
C            10.12.30.0/24 is directly connected, eth0, 02:01:45
C            20.20.20.0/24 is directly connected, eth2, 01:02:04
C            30.30.30.0/24 is directly connected, eth3, 01:02:04
i L1         40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:35:32
i L1         50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:35:45
i L1         60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:35:32
C            127.0.0.0/8 is directly connected, lo, 02:01:45

Gateway of last resort is not set

R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  40.40.40.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 25
    Protection Provided : LP NP BP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info
       * - candidate default

IP Route Table for VRF "default"
i L1    40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:35:48
                      [FRR-NH] via 30.30.30.144, eth3  FRR Table has only
allowed prefix through route-map , for remaining prefixes, FRR not present
```

LFA Tie-Breaker

Based on the index values configured, if inequalities are satisfied, protections will be provided:

- Lower the index will have the highest priority, The path which provides protection with highest priority will be selected. If there are multiple paths providing the highest priority protection then we will check which of the path provides the protection which has 2nd highest priority and so on.
- If all the paths provide same priority, then the LFA route is chosen on the basis of path cost.
- If none of the paths provides the protection with highest priority, then we will see which path provides the 2nd highest priority and so on.

The show command below displays default values for tie-breaker, by default maximum protection (link, node, broadcast, if ecmp, ecmp backup path) will be provided.

After configuring tie-breaker with index, values will be changed accordingly.

```
R1#show ip isis lfa-config level-1

TIE-Breaker Preference values
-----------------------------------------------
 Primary Path                 : 20
 Link Protecting              : 60
 Node Protecting              : 30
 Broadcast Interface Disjoint : 70
 Secondary Path               : 255
 Downstream Path              : 90

 Termination Hold On Interval : 1000 ms
```

R1

To change index values, below configurations should be used, with the lower the index highest the priority.

`(config)#router isis 1`	Create an IS-IS routing instance for area `49` with instance 1
`(config-router)#fast-reroute per-prefix level-1 proto ipv4 all`	Configure LFA-FRR to calculate the available backup path for all L1 ipv4 prefixes learned
`(config-router)#fast-reroute tie-break level-1 proto ipv4 interface-disjoint index 1`	Configure index value to change priority for link protection
`(config-router)#fast-reroute tie-break level-1 proto ipv4 node-protecting index 2`	Configure index value to change priority for node protection
`(config-router)#fast-reroute tie-break level-1 proto ipv4 broadcast-interface-disjoint index 3`	Configure index value to change priority for broadcast link protection
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

```
R1#show ip isis lfa-config level-1

TIE-Breaker Preference values
-----------------------------------------------
 Primary Path                 : 20
 Link Protecting              : 1
 Node Protecting              : 2
 Broadcast Interface Disjoint : 3
 Secondary Path               : 255
 Downstream Path              : 90

 Termination Hold On Interval : 1000 ms
```

```
R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  40.40.40.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 25
    Protection Provided : LP NP BP

L1  50.50.50.0/24
    Primary Path via    : 20.20.20.143, eth2
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 20
    Protection Provided : LP NP BP

L1  60.60.60.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 20
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info
       * - candidate default

IP Route Table for VRF "default"
i L1    40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:45:16
                      [FRR-NH] via 30.30.30.144, eth3

i L1    50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:45:29
                      [FRR-NH] via 30.30.30.144, eth3

i L1    60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:45:16
                      [FRR-NH] via 30.30.30.144, eth3
```

LFA Termination

A router MUST limit the amount of time an alternate next-hop is used after the primary next-hop has become unavailable. This ensures that the router will start using the new primary next-hops.

LFA termination avoids a micro looping in topology, when particular network goes down, LFA backup path will be installed and if termination interval is configured, LFA backup will be still used till the interval and it is used in order to verify new primary path is loop free.

R1

Configure termination interval on R1 in router mode:

`(config)#router isis 1`	Create an IS-IS routing instance for area `49` with instance 1
`(config-router)#fast-reroute terminate-hold-on interval 100000`	Configure LFA termination interval
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

```
R1#show ip isis lfa-config level-1

TIE-Breaker Preference values
-------------------------------------------------
 Primary Path                 : 20
 Link Protecting              : 1
 Node Protecting              : 2
 Broadcast Interface Disjoint : 3
 Secondary Path               : 255
 Downstream Path              : 90

 Termination Hold On Interval : 100000 ms

R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth1, 01:19:46
C            10.12.30.0/24 is directly connected, eth0, 02:19:27
C            20.20.20.0/24 is directly connected, eth2, 01:19:46
C            30.30.30.0/24 is directly connected, eth3, 01:19:46
i L1         40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:53:14
i L1         50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:53:27
i L1         60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:53:14
C            127.0.0.0/8 is directly connected, lo, 02:19:27

Gateway of last resort is not set
```

Shut down one of the primary nexthops, here eth2 of R1:

`(config)#interface eth2`	Enter interface mode
`(config-if)#shutdown`	Shutdown the interface

(config-if)#exit	Exit interface mode
(config)exit	Exit config mode

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.002a.230a  Up     27        L1   IS-IS
0000.0000.0004 eth3        5254.00f5.35a4  Up     7         L1   IS-IS
```

Here, eth1 has become a primary path, which was originally a backup path:

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth1, 01:24:47
C            10.12.30.0/24 is directly connected, eth0, 02:24:28
C            30.30.30.0/24 is directly connected, eth3, 01:24:47
i L1         40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:02:01
i L1         50.50.50.0/24 [115/20] via 10.10.10.141, eth1, 00:02:01
i L1         60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:02:01
C            127.0.0.0/8 is directly connected, lo, 02:24:28

Gateway of last resort is not set

R1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
C    10.10.10.0/24       5           --                  eth1            0
C    30.30.30.0/24       15          --                  eth3            0
L1   40.40.40.0/24       10          10.10.10.141        eth1            0
L1   50.50.50.0/24       20          10.10.10.141        eth1            0
L1   60.60.60.0/24       15          10.10.10.141        eth1            0

R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
```

```
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  40.40.40.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 25
    Protection Provided : LP NP BP

L1  50.50.50.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 30
    Protection Provided : LP NP BP DP

L1  60.60.60.0/24
    Primary Path via    : 10.10.10.141, eth1
    FRR Backup Path via : 30.30.30.144, eth3
    FRR Metric          : 20
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info
       * - candidate default

IP Route Table for VRF "default"

i L1    40.40.40.0/24 [115/10] via 10.10.10.141, eth1, 00:02:19
                      [FRR-NH] via 30.30.30.144, eth3

i L1    50.50.50.0/24 [115/20] via 10.10.10.141, eth1, 00:02:19
                      [FRR-NH] via 30.30.30.144, eth3

i L1    60.60.60.0/24 [115/15] via 10.10.10.141, eth1, 00:02:19
                      [FRR-NH] via 30.30.30.144, eth3
```

LFA For ECMP Paths

Equal-Cost Multi-Path Routing (ECMP) is a routing technique where next-hop packet forwarding to a single destination can occur over multiple “best-paths” that tie in routing metric calculations. Because it is a per-hop decision limited to a single router, it can increase bandwidth by load-balancing traffic over multiple paths.

Here, we provide configuration capabilities for Loop-Free Alternate (LFA) Fast Reroute (FRR) along with ECMP.

Topology

Figure 1-16 shows the configuration to enable the ISIS LFA feature with ECMP.

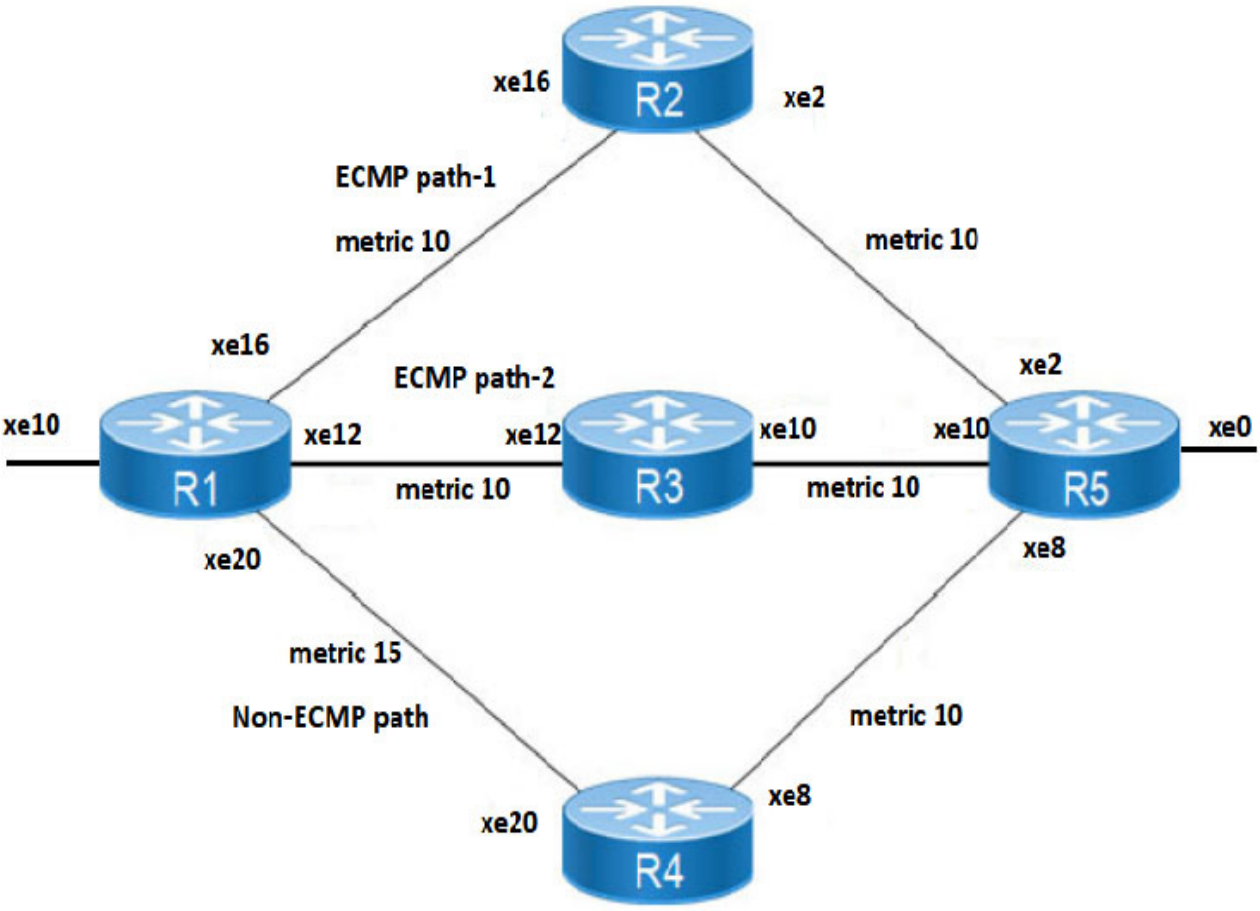

Figure 1-16: ISIS LFA-FRR ECMP

R1

`#configure terminal`	Enter configure mode.
`(config)#int xe16`	Enter interface mode.
`(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe12`	Enter interface mode.
`(config-if)#ip address 20.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe20`	Enter interface mode.
`(config-if)#ip address 30.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 15`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe10`	Enter interface mode.
`(config-if)#ip address 31.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 with instance 1
`(config-router)# net 49.0001.0000.0000.0001.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-router)#fast-reroute per-prefix level-1 proto ipv4 all`	Configure LFA-FRR to calculate the available backup path for all L1 ipv4 prefixes learnt
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#int xe16`	Enter interface mode.
`(config-if)#ip address 10.1.1.2/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe2`	Enter interface mode.
`(config-if)#ip address 40.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 with instance 1
`(config-router)# net 49.0001.0000.0000.0002.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R3

`#configure terminal`	Enter configure mode.
`(config)#int xe12`	Enter interface mode.
`(config-if)#ip address 20.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe10`	Enter interface mode.
`(config-if)#ip address 50.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 with instance 1
`(config-router)# net 49.0001.0000.0000.0003.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R4

`#configure terminal`	Enter configure mode.
`(config)#int xe20`	Enter interface mode.
`(config-if)#ip address 30.1.1.2/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1g`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#int xe8`	Enter interface mode.
`(config-if)#ip address 60.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip router isis 1`	Enable ISIS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure isis metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 with instance 1
`(config-router)# net 49.0001.0000.0000.0004.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

RTR5

#configure terminal	Enter configure mode.
(config)#int xe2	Enter interface mode.
(config-if)#ip address 40.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 10	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#int xe10	Enter interface mode.
(config-if)#ip address 50.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 10	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#int xe8	Enter interface mode.
(config-if)#ip address 60.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 10	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#int xe0	Enter interface mode.
(config-if)#ip address 21.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip router isis 1	Enable ISIS routing on interface for area 49 with instance 1
(config-if)#isis metric 10	Configure isis metric value for interface
(config-if)#isis circuit-type level-1	Enable circuit type on interface
(config-if)#exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 with instance 1
(config-router)# net 49.0001.0000.0000.0005.00	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for ISIS on all interfaces
(config-router)#exit	Exit router mode.
(config)#exit	Exit config mode.

Validation

```
R1 (Source):
The backup path will be selected by default from same Primary/ECMP set and
"Pri" indicates backup selected from ECMP set.
R1#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
```

```
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0003 xe12        e8c5.7a6b.732a  Up     21        L1   IS-IS
0000.0000.0002 xe16        e8c5.7a7d.532e  Up     21        L1   IS-IS
0000.0000.0004 xe20        e8c5.7a25.2752  Up     19        L1   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0005 xe2         e8c5.7a76.5820  Up     28        L1   IS-IS
0000.0000.0001 xe16        e8c5.7af5.ef2e  Up     7         L1   IS-IS

R3#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0005 xe10        e8c5.7a76.5828  Up     8         L1   IS-IS
0000.0000.0001 xe12        e8c5.7af5.ef2a  Up     5         L1   IS-IS

R4#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0005 xe8         e8c5.7a76.5826  Up     8         L1   IS-IS
0000.0000.0001 xe20        e8c5.7af5.ef32  Up     8         L1   IS-IS

R5#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 xe2         e8c5.7a7d.5320  Up     8         L1   IS-IS
0000.0000.0004 xe8         e8c5.7a25.2746  Up     29        L1   IS-IS
0000.0000.0003 xe10        e8c5.7a6b.7328  Up     26        L1   IS-IS

R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  21.1.1.0/24
```

```
    Primary Path via    : 10.1.1.2, xe16
    FRR Backup Path via : 20.1.1.2, xe12
    FRR Metric          : 30
    Protection Provided : LP NP BP Pri DP  >>> Here Pri indicates backup
selected from ECMP set

    Primary Path via    : 20.1.1.2, xe12
    FRR Backup Path via : 10.1.1.2, xe16
    FRR Metric          : 30
    Protection Provided : LP NP BP Pri DP

L1  40.1.1.0/24
    Primary Path via    : 10.1.1.2, xe16
    FRR Backup Path via : 20.1.1.2, xe12
    FRR Metric          : 30
    Protection Provided : LP NP BP

L1  50.1.1.0/24
    Primary Path via    : 20.1.1.2, xe12
    FRR Backup Path via : 10.1.1.2, xe16
    FRR Metric          : 30
    Protection Provided : LP NP BP

L1  60.1.1.0/24
    Primary Path via    : 30.1.1.2, xe20
    FRR Backup Path via : 20.1.1.2, xe12
    FRR Metric          : 30
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info, E - EVPN
       * - candidate default

IP Route Table for VRF "default"
i L1    21.1.1.0/24 [115/30] via 20.1.1.2, xe12, 00:00:57
                    [FRR-NH] via 10.1.1.2, xe16

                   [115/30] via 10.1.1.2, xe16
                    [FRR-NH] via 20.1.1.2, xe12

i L1    40.1.1.0/24 [115/20] via 10.1.1.2, xe16, 00:07:52
                    [FRR-NH] via 20.1.1.2, xe12

i L1    50.1.1.0/24 [115/20] via 20.1.1.2, xe12, 00:07:12
                    [FRR-NH] via 10.1.1.2, xe16

i L1    60.1.1.0/24 [115/25] via 30.1.1.2, xe20, 00:06:44
                    [FRR-NH] via 20.1.1.2, xe12

R1#show ip isis lfa-config level-1

TIE-Breaker Preference values
```

```
-------------------------------------------------
 Primary Path                 : 20
 Link Protecting              : 60
 Node Protecting              : 30
 Broadcast Interface Disjoint : 70
 Secondary Path               : 255
 Downstream Path              : 90

 Termination Hold On Interval : 1000 ms
```

Backup Path for ECMP Path from Non-ECMP Path

To select Backup path from secondary/Non-ECMP path, configure the below command in R1 with lowest index value.

If no backup path available from non-ecmp set , then backup path is installed from primary set.

`(config)#router isis 1`	Create an IS-IS routing instance for area `49` with instance 1
`(config-router)#net 49.0000.0000.0001.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#fast-reroute tie-break level-1 proto ipv4 secondary-path index 1`	Configure secondary path tie-breaker to select backup path from Non-ECMP path
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

```
R1#show ip isis lfa-config level-1

TIE-Breaker Preference values
-------------------------------------------------
 Primary Path                 : 20
 Link Protecting              : 60
 Node Protecting              : 30
 Broadcast Interface Disjoint : 70
 Secondary Path               : 1
 Downstream Path              : 90

 Termination Hold On Interval : 1000 ms

R1#Below, "Sec" indicates the backup path is from the Non-ECMP path
R1#show ip isis route fast-reroute

Tag   : 1  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
        D - discard, LP - Link Protecting, NP - Node Protecting,
        BP - Broadcast Interface Disjoint, Pri - Primary Path,
        Sec - Secondary Path, DP - Downstream Path

L1  21.1.1.0/24
    Primary Path via    : 10.1.1.2, xe16
    FRR Backup Path via : 30.1.1.2, xe20
    FRR Metric          : 35
```

```
    Protection Provided : LP NP BP Sec DP >>> Sec indicates backup path is
from Non-ecmp path

    Primary Path via    : 20.1.1.2, xe12
    FRR Backup Path via : 30.1.1.2, xe20
    FRR Metric          : 35
    Protection Provided : LP NP BP Sec DP

L1  40.1.1.0/24
    Primary Path via    : 10.1.1.2, xe16
    FRR Backup Path via : 20.1.1.2, xe12
    FRR Metric          : 30
    Protection Provided : LP NP BP

L1  50.1.1.0/24
    Primary Path via    : 20.1.1.2, xe12
    FRR Backup Path via : 10.1.1.2, xe16
    FRR Metric          : 30
    Protection Provided : LP NP BP

L1  60.1.1.0/24
    Primary Path via    : 30.1.1.2, xe20
    FRR Backup Path via : 20.1.1.2, xe12
    FRR Metric          : 30
    Protection Provided : LP NP BP DP

R1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area ,p - stale info, E - EVPN
       * - candidate default

IP Route Table for VRF "default"
i L1    21.1.1.0/24 [115/30] via 20.1.1.2, xe12, 00:02:48
                    [FRR-NH] via 30.1.1.2, xe20

                   [115/30] via 10.1.1.2, xe16
                    [FRR-NH] via 30.1.1.2, xe20

i L1    40.1.1.0/24 [115/20] via 10.1.1.2, xe16, 00:09:43
                    [FRR-NH] via 20.1.1.2, xe12

i L1    50.1.1.0/24 [115/20] via 20.1.1.2, xe12, 00:09:03
                    [FRR-NH] via 10.1.1.2, xe16

i L1    60.1.1.0/24 [115/25] via 30.1.1.2, xe20, 00:08:35
                    [FRR-NH] via 20.1.1.2, xe12
```

CHAPTER 2 IS-IS IPv6 Configuration

This chapter contains basic IS-IS (Intermediate System to Intermediate System) on IPv6 configuration examples.

Enable IS-ISv6 on an Interface

This example shows the minimum configuration required for enabling IS-IS on IPv6 on an interface. R1 and R2 are two routers in the ABC instance connecting to the network 1000::/64. After enabling IS-IS on an interface, create a routing instance, and specify the Network Entity Title (NET). IS-IS explicitly specifies a NET to begin routing. NET is comprised of the area address and the system ID of the router.

Note: ISISv6 session will come up even if IPv6 address is not configured, as it will use the link local address present on the interfaces.

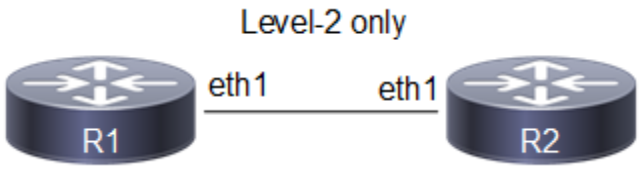

Figure 2-17: Figure 4-46: Basic IS-IS v6 Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).

© 2025 IP Infusion Inc. Proprietary

`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5      Up     27        L2   IS-IS

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     7         L2   IS-IS
R2#

R1#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
C    1000::/64 [10]
      via ::, eth1

R1#

R2#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
C    1000::/64 [10]
      via ::, eth1

R2#

R1#
R1#show ipv6 isis topology
```

```
Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       --
0000.0000.0002       10         0000.0000.0002       eth1        b86a.97c4.31c5

R1#

R2#show ipv6 isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
0000.0000.0001       10         0000.0000.0001       eth1        b86a.97cb.3ec5
0000.0000.0002       --
```

Set Priority

This example describes how to set the priority for an interface. Set a high priority for a router to make it the Designated IS (DIS). Router R3 is configured to have a priority of 70, this is higher than the default priority (64) of R1 and R2. This makes R3 the DIS.

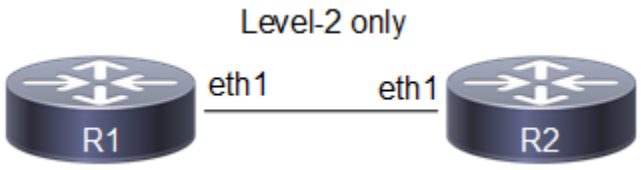

Figure 2-18: Set IS-IS Priority

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address on interface.
`(config-if)#isis priority 125`	Specify the router priority to a higher priority (125) to make R2 the designated IS (DIS).
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5      Up     8         L2   IS-IS
R1#

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     26        L2   IS-IS
R2#


R1#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0002 eth1        Up     L2   125       0000.0000.0002.01
R1#

R2#show clns is-neighbors

Tag ABC:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
```

```
0000.0000.0001 eth1        Up     L2   64        0000.0000.0002.01
R2#


R1#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00002722
    Local SNPA: b86a.97cb.3ec5
    IP interface address:
    IPv6 interface address:
      1000::1/64
      fe80::ba6a:97ff:fecb:3ec5/64
    Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0002.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 0 milliseconds
R1#

R2#show isis interface
eth1 is up, line protocol is up
  Routing Protocol: IS-IS (ABC)
    Network Type: Broadcast
    Circuit Type: level-1-2
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00002722
    Local SNPA: b86a.97c4.31c5
    IP interface address:
    IPv6 interface address:
      1000::2/64
      fe80::ba6a:97ff:fec4:31c5/64
    Level-2 Metric: 10/10, Priority: 125, Circuit ID: 0000.0000.0002.01
    Number of active level-2 adjacencies: 1
    Level-2 LSP MTU: 1492
    Next IS-IS LAN Level-2 Hello in 1 seconds
R2#
```

Dynamic hostname

This example shows how to configure Dynamic Hostname for an ISIS IPv6 instance. Dynamic hostname is the method of mapping name-to-systemID. It allows the routing protocol to advertise symbolic names in the IS-IS PDUs. This is done by the addition of a new TLV which allows the IS-IS routers to include the name-to-systemID mapping data in their LSPs. This allows for simple and reliable transport of name mapping across IS-IS networks.

Dynamic hostname can be either the hostname of the node or the tag of the configured ISISv6 instance.

Note: Dynamic-hostname has to be configured on all nodes for it to take effect.

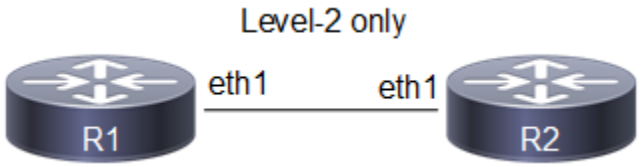

Figure 2-19: Basic dynamic hostname topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#dynamic-hostname`	Configure the hostname to be advertised for an ISIS instance.
`(config-router)#end`	Exit the current mode and enter privilege mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#dynamic-hostname`	Configure the hostname to be advertised for an ISIS instance.
`(config-router)#end`	Exit the current mode and enter privilege mode.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
```

```
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R2             eth1        b86a.97c4.31c5      Up     20        L2   IS-IS
R1#

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             eth1        b86a.97cb.3ec5      Up     8         L2   IS-IS
R2#



R1#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000008   0xFB86        1144              0/0/0
R1.01-00            * 0x00000005   0x19BD        1141              0/0/0
R2.00-00              0x00000007   0x245C        1140              0/0/0

R1#

R2#show isis database
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000008   0xFB86        1144              0/0/0
R1.01-00              0x00000005   0x19BD        1140              0/0/0
R2.00-00            * 0x00000007   0x245C        1140              0/0/0
R2.01-00            * 0x00000002   0xE710        0 (1132)          0/0/0

R2#



R1#show ipv6 isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       eth1        b86a.97c4.31c5

R1#
```

```
R2#show ipv6 isis topology

Tag ABC:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   10         R1       eth1        b86a.97cb.3ec5
R2                   --

R2#
```

Redistribute Routes into IS-IS

In this example, the configuration causes OSPFv3 routes to be imported into the IS-ISv6 routing table, and advertised into the ABC instance.

Topology

Figure 2-20: Redistribute Routes into IS-IS

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.

(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on an interface and assign the Area ID 0.
(config-if)#ipv6 address 2000::1/64	Configure IPv6 address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0005.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#address-family ipv6	Enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration
(config-router-af)#redistribute ospf	Enable redistribution of routes from ospf into the ISIS routing table.
(config-router-af)#exit-address-family	Exit address family mode.
(config-router)#exit	Exit router mode.
(config)#router ipv6 ospf	Create an OSPFv3 routing instance.
(config-router)#router-id 2.2.2.2	Specify a Router ID (2.2.2.2) for the OSPFv3 routing process.
(config-router)#end	Exit the current mode and enter privilege mode.

R3

#configure terminal	Enter configure mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on an interface and assign the Area ID 0.
(config-if)#ipv6 address 2000::2/64	Configure IPv6 address on interface.
(config-if)#exit	Exit interface mode.
(config)#router ipv6 ospf	Create an OSPFv3 routing instance.
(config-router)#router-id 3.3.3.3	Specify a Router ID (3.3.3.3) for the OSPFv3 routing process.
(config-router)#end	Exit the current mode and enter privilege mode.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5  Up     27        L2   IS-IS
R1#

R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
```

```
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     8         L2   IS-IS
R2#

R2#show ipv6 ospf neighbor

Total number of full neighbors: 1
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
3.3.3.3           1   Full/DR         00:00:34    eth2        0
R2#


R1#show ipv6 route isis
IP Route Table for VRF "default"
i L2   2000::/64 [115/10] via fe80::ba6a:97ff:fec4:31c5, eth1, 00:21:19
R1#


R1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:33:07
C      1000::/64 via ::, eth1, 01:13:36
i L2   2000::/64 [115/10] via fe80::ba6a:97ff:fec4:31c5, eth1, 00:21:29
C      fe80::/64 via ::, xe8, 00:39:44
R1#
```

Interface Metric

You can make a route the preferred route by changing its metric. In this example, the cost has been configured to make R3 the next hop for R1.

The default metric for each interface is 10. Interface eth2 on R2 has a metric of 20, and Interface eth2 on R3 has a metric of 30. The total cost to reach 9999::/64 (R4) through R2 and R3 is computed as follows: R2: 10+20 = 30 R3: 10+30 = 40

In this topology, R1 chooses R2 as its next hop for destination 9999::/64.

Note: Below configuration is applicable for narrow (non-wide) metric-style. Use the "metric-style wide" under isis instance and "isis wide-metric < 1-16777214>" under interface mode CLI’s to configure Wide metric.

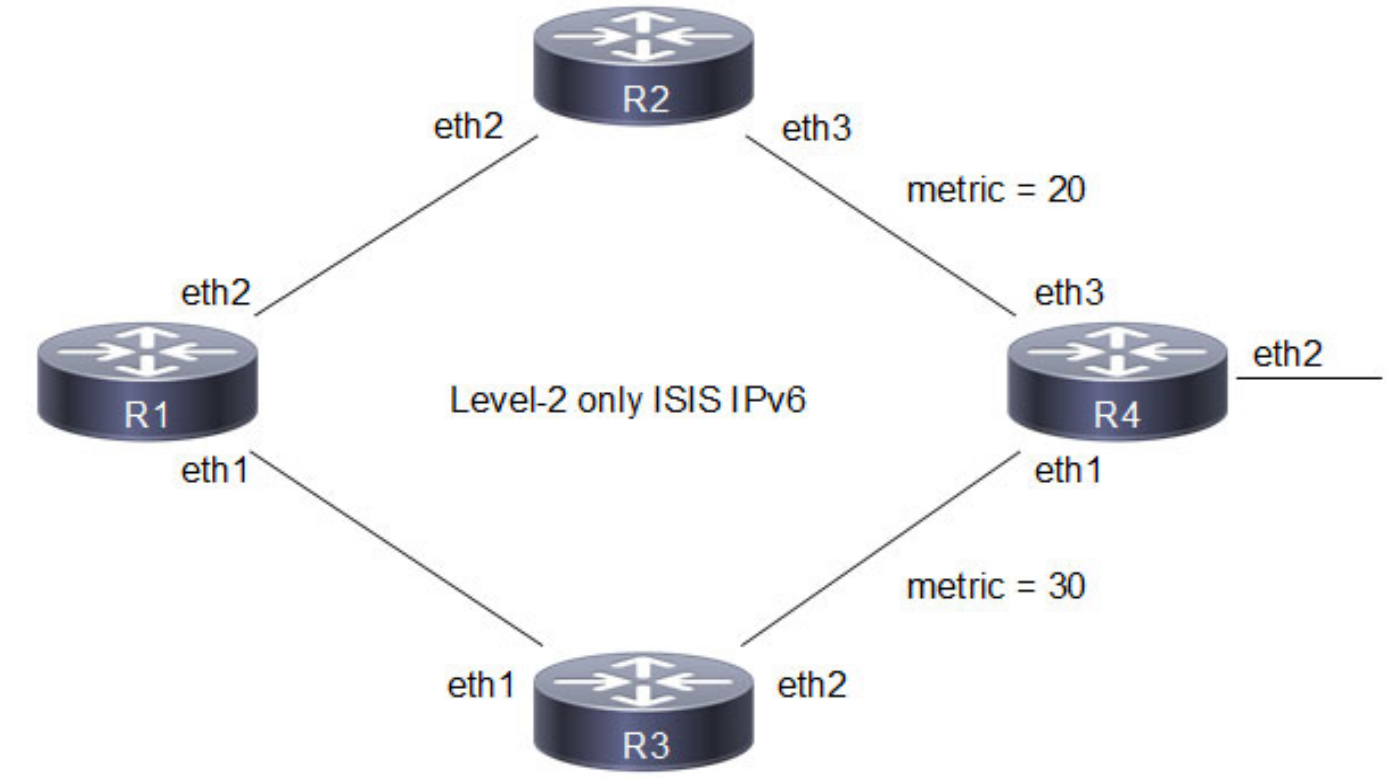

gy

Figure 2-21: Configure IS-IS Metric

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-IS routing on an interface for area 49(ABC).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-IS routing on an interface for area 49(ABC).
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0100.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#end`	Exit current mode and enter privilege mode.

R2

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).

(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis metric 20	Set the value of IS-IS metric (on eth2) to 20.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0100.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R3

(config)#interface eth1	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#isis metric 30	Set the value of IS-IS metric (on eth2) to 30.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0100.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R4

(config)#interface eth1	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 address 9999::1/64	Configure ipv6 address in eth2
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-IS routing on an interface for area 49 (ABC).
(config-if)#ipv6 address 2000::2/64	Configure IPv6 address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0100.0000.0000.0004.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#address-family ipv6	Enter ipv6 address family.
(config-router-af)#redistribute connected	Enable redistribution of connected routes into ISIS process
(config-router-af)#end	Exit current mode and enter privilege mode

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        b86a.97c4.31c5      Up     25        L2   IS-IS
0000.0000.0003 eth1         b86a.97c9.3cc5      Up     26        L2   IS-IS
R1#


R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0004 eth3         b86a.97c7.32c5      Up     7         L2   IS-IS
0000.0000.0001 eth2        b86a.97cb.3ec5      Up     6         L2   IS-IS
R2#

R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0004 eth2        b86a.97c7.32c5      Up     22        L2   IS-IS
0000.0000.0001 eth1         b86a.97cb.3ec5      Up     7         L2   IS-IS
R3#

R4#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 2
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth3         b86a.97c4.31c5      Up     22        L2   IS-IS
0000.0000.0003 eth1        b86a.97c9.3cc5      Up     7         L2   IS-IS
R4#

R1#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
C    1000::/64 [10]
      via ::, ce0
L2   2000::/64 [30]
      via fe80::ce37:abff:fe87:3a74, ce0
C    3000::/64 [10]
      via ::, xe14
L2   4000::/64 [40]
      via fe80::ba6a:97ff:fecf:3ad4, xe14
      via fe80::ce37:abff:fe87:3a74, ce0
L2   9999::/64 [30]
      via fe80::ce37:abff:fe87:3a74, ce0
R1#
```

Route Summarization

Route summarization makes the routing table smaller, but still allows complete IP connectivity.

The following example consists of a three-router topology, in which R2 is doing the summarization. In this example, R1 is the L1 router, R2 is the L1/L2 router doing the summarization, and R3 is the L2 router. The following configuration is given only for R2, assuming that the adjacencies with R1 and R3 are already up, and the route tables with the appropriate routes are already populated.

Topology

Figure 2-22: Route Summarization Topology

Configuration

R1

h

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000 ::1/64`	Configure IPv6 address on interface.
`(config-if)#isis circuit-type level-1`	Set the circuit type as level-1 for the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).

(config-if)#ipv6 address 1000 ::1/64	Configure IPv6 address on interface.
(config-router)#is-type level-1	Configure instance as level-1 routing.
(config-router)#net 49.0001.0000.0000.0001.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

R2

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-ISv6 routing on an interface for area 49(ABC).
(config-if)#ipv6 address 1000 ::2/64	Configure IPv6 address on interface.
(config-if)#isis circuit-type level-1	Set the circuit type as level-1 for the interface.
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-ISv6 routing on an interface for area 49(ABC).
(config-if)#ipv6 address 2000 ::1/64	Configure IPv6 address on interface.
(config-if)#isis circuit-type level-2-only	Set the circuit type as level-2-only for the interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#net 49.0001.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#address-family ipv6	Enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration.
(config-router-af)#redistribute isis level-2 into level-1	Enable redistribution of isis routes from level-2 into level-1
(config-router-af)#summary-prefix 11:1:1:1::/64 level-1 metric 58	Configure the summary prefix to summarize IPv6 reachability information.
(config-router-af)#exit-address-family	Exit address family mode.
(config-router)#exit	Exit router mode.

R3

#configure terminal	Enter configure mode.
(config)#ipv6 route 11:1:1:1:1:1::/96 eth2	Configure ipv6 static route.
(config)#ipv6 route 11:1:1:1:2:1::/96 eth2	Configure ipv6 static route.
(config)#ipv6 route 11:1:1:1:3:1::/96 eth2	Configure ipv6 static route.
(config)#interface ethh2	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-ISv6 routing on an interface for area 49(ABC).
(config-if)#ipv6 address 2000 ::2/64	Configure IPv6 address on interface.
(config-if)#isis circuit-type level-2-only	Set the circuit type as level-2-only for the interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#net 49.0001.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.

(config-router)#is-type level-2-only	Configure instance as level-2 -onlyrouting.
(config-router)#address-family ipv6	Enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration.
(config-router-af)#redistribute static	Enable redistribution of static routes into ISIS instance.
(config-router-af)#end	Exit the current mode and enter privilege mode.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5      Up     21        L1   IS-IS
R1#


R2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0003 eth2         b86a.97c7.32c5      Up     8         L2   IS-IS
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     8         L1   IS-IS
R2#


R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2         b86a.97c4.31c5      Up     20        L2   IS-IS
R3#



R1#show ipv6 route isis
IP Route Table for VRF "default"
i ia   11:1:1:1::/64 [115/68] via fe80::ba6a:97ff:fec4:31c5, eth1, 00:02:29
R1#

R2#show ipv6 route isis
```

```
IP Route Table for VRF "default"
i      11:1:1:1::/64 [115/0] via ::, Null, 00:02:30
i L2   11:1:1:1:1:1::/96 [115/10] via fe80::ba6a:97ff:fec7:32c5, eth2, 00:04:54
i L2   11:1:1:1:2:1::/96 [115/10] via fe80::ba6a:97ff:fec7:32c5, eth2, 00:04:54
i L2   11:1:1:1:3:1::/96 [115/10] via fe80::ba6a:97ff:fec7:32c5, eth2, 00:04:54
R2#

R3#show ipv6 route isis
IP Route Table for VRF "default"
R3#
```

```
R1#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
ia   11:1:1:1::/64 [68]
      via fe80::ba6a:97ff:fec4:31c5, eth1

R1#
```

```
R2#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
D    11:1:1:1::/64 [0]
      via ::
L2   11:1:1:1:1:1::/96 [10]
      via fe80::ba6a:97ff:fec7:32c5, eth2
L2   11:1:1:1:2:1::/96 [10]
      via fe80::ba6a:97ff:fec7:32c5, eth2
L2   11:1:1:1:3:1::/96 [10]
      via fe80::ba6a:97ff:fec7:32c5, eth2

R2#
```

```
R3#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
```

```
E    11:1:1:1:1:1::/96 [0]
      via ::
E    11:1:1:1:2:1::/96 [0]
      via ::
E    11:1:1:1:3:1::/96 [0]
      via ::

R3#



R1#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x00000002   0xA557        497               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0001.01
0000.0000.0001.01-00* 0x00000001   0x1FBD        497               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00  0x00000004   0x4DA0        909               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0001.01
  Metric:   58         IPv6-Interarea 11:1:1:1::/64

R1#


R2#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000002   0xA557        496               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0001.01
0000.0000.0001.01-00  0x00000001   0x1FBD        496               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000004   0x4DA0        910               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0001.01
  Metric:   58         IPv6-Interarea 11:1:1:1::/64

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
```

```
0000.0000.0002.00-00* 0x00000002   0xCD2A        506               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0003.01
0000.0000.0003.00-00  0x00000003   0x89E2        765               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0003.01
  Metric:   0          IPv6 11:1:1:1:1:1::/96
  Metric:   0          IPv6 11:1:1:1:2:1::/96
  Metric:   0          IPv6 11:1:1:1:3:1::/96
0000.0000.0003.01-00  0x00000001   0x2DA9        505               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00

R2#
```

Passive Interface

In ISP and large enterprise networks, many of the distribution routers have more than 200 interfaces. Before the Default Passive-Interface feature, there were two possibilities for obtaining routing information from tall of these interfaces:

- Configure a routing protocol on the backbone interfaces and redistribute connected interfaces.
- Configure the routing protocol on all interfaces and manually set most of them as passive, which was time consuming.

The solution to this problem was to configure the routing protocol on all interfaces and manually set the passive interface command on the interfaces where adjacency was not desired. In certain networks, this meant coding 200 or more passive-interface statements. With the Default Passive Interface feature, this problem is solved by allowing all interfaces to be set as passive by default using a single passive-interface default command, then configuring individual interfaces in which adjacencies are desired using the no passive-interface command.

Usage

1. When a specific interface is configured as passive using the passive-interface command:
   - The interface loses its adjacency on that interface, for example, eth1.
   - The interface (eth1) is still advertised by other IS-IS speaking interfaces to their neighbors.
2. When a specific interface is configured as passive using passive-interface command followed by removing the configuration using no passive-interface command:
   - The interface is IS-IS disabled and must be enabled using the ipv6 router isis command (for example, ipv6 router isis 1).
   - The interface (for example, eth1) is not advertised by other IS-IS speaking interfaces to their neighbors.
3. When an interface is configured with the passive interface command:
   - All interfaces lose their adjacency, except the interface with the higher index number. (For example: If eth1, eth2, eth3, and eth4 are the router interfaces, eth4 has the highest index number.)
   - All interfaces are advertised by the active IS-IS speaking interface to its neighbors.

4. When an interface is configured with the no passive interface command:
   - All interfaces are IS-IS disabled, except the interface that was active, and all interfaces must enable IS-IS on these interfaces using the ipv6 router isis command (for example, ipv6 router isis 1).
   - All interfaces are not advertised by the active IS-IS speaking interface to its neighbors.

Topology

Figure 2-23: IS-ISv6 Passive Interface

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#end`	Exit the current mode and enter privilege mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#ipv6 address 1000::2/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable OSPFv3 routing on an interface and assign the Area ID 0.
`(config-if)#ipv6 address 2000::1/64`	Configure IPv6 address on interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.

(config-router)#net 49.0005.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#passive-interface eth1	Configure interface eth1 as a passive-interface.
(config-router)#end	Exit the current mode and enter privilege mode.

R3

#configure terminal	Enter configure mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 router isis ABC	Enable IS-ISv6 routing on an interface for area 49(ABC).
(config-if)#ipv6 address 2000::2/64	Configure IPv6 address on interface.
(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0005.0000.0000.0003.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#end	Exit the current mode and enter privilege mode.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 0
Total number of adjacencies: 0
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1#


R2#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0003 eth2         b86a.97c7.32c5   Up     7         L2   IS-IS
R2#


R3#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
```

```
0000.0000.0002 eth2         b86a.97c4.31c5      Up     24         L2   IS-IS
R3#


R1#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
C    1000::/64 [10]
      via ::, eth1

R1#


R2#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
C    1000::/64 [0]
      via ::, eth1
C    2000::/64 [10]
      via ::, eth2

R2#

R3#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
L2   1000::/64 [10]
      via fe80::ba6a:97ff:fec4:31c5, eth2
C    2000::/64 [10]
      via ::, eth2

R3#


R1#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00* 0x0000000E   0x0E19        931               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
```

```
  IPv6 Address: 1000::1
  Metric:   10         IPv6 1000::/64
0000.0000.0001.01-00* 0x00000009   0xE110        0 (931)           0/0/0
0000.0000.0002.00-00  0x00000011   0xABC7        440               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 1000::2
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 1000::/64
  Metric:   10         IPv6 2000::/64
0000.0000.0003.00-00  0x00000002   0x9ED3        439               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 2000::2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
0000.0000.0003.01-00  0x00000001   0x2DA9        436               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00

R1#
R1#

R2#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x0000000D   0x2666        879               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 1000::1
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IPv6 1000::/64
0000.0000.0001.01-00  0x00000009   0x11C1        821               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000012   0x669F        906               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 2000::1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   0          IPv6 1000::/64
  Metric:   10         IPv6 2000::/64
0000.0000.0003.00-00  0x00000002   0x9ED3        439               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 2000::2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
```

```
0000.0000.0003.01-00  0x00000001   0x2DA9        436               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00

R2#

R3#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x0000000D   0x2666        879               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 1000::1
  Metric:   10         IS 0000.0000.0001.01
  Metric:   10         IPv6 1000::/64
0000.0000.0001.01-00  0x00000009   0x11C1        821               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS 0000.0000.0002.00
0000.0000.0002.00-00  0x00000012   0x669F        905               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 2000::1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   0          IPv6 1000::/64
  Metric:   10         IPv6 2000::/64
0000.0000.0003.00-00* 0x00000002   0x9ED3        440               0/0/0
  Area Address: 49.0005
  NLPID:        0x8E
  IPv6 Address: 2000::2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
0000.0000.0003.01-00* 0x00000001   0x2DA9        437               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00

R3#
```

Enable BFD over IS-ISv6

This example shows how to configure Bidirectional Forwarding Detection with ISISv6 instance

Topology

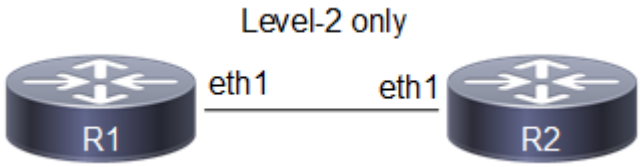

Figure 2-24: Basic BFD over IS-ISv6 Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance.
`(config-router)#end`	Exit current mode and enter privilege mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0005.0000.0000.0002.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance.
`(config-router)#end`	Exit current mode and enter privilege mode.

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
```

```
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5  Up     24        L2   IS-IS
R1#

R1#show bfd session

BFD process for VRF: (DEFAULT VRF)
------------------------------------------------------------------------------
Sess-Idx   Remote-Disc  Lower-Layer  Sess-Type   Sess-State  UP-Time   Interface
Down-Reason  Remote-Addr
256      256        IPv6       Single-Hop  Up        00:04:26  eth1         NA
fe80::ba6a:97ff:fec4:31c5/128
Number of Sessions:    1
R1#
R1#show bfd session detail

BFD process for VRF: (DEFAULT VRF)
------------------------------------------------------------------------------

Session Interface Index : 10018          Interface name :eth1
Session Index : 256
Lower Layer : IPv6                       Version : 1
Session Type : Single Hop                Session State : Up
Local Discriminator : 256                Local Address : fe80::ba6a:97ff:fecb:3ec5/128
Remote Discriminator : 256               Remote Address : fe80::ba6a:97ff:fec4:31c5/128
Local Port : 49152                       Remote Port : 3784
Options :

Diagnostics : None

Timers in Milliseconds :
Min Tx: 250              Min Rx: 250              Multiplier: 3
Neg Tx: 250              Neg Rx: 250              Neg detect mult: 3
Min echo Tx: 1000        Min echo Rx: 1000        Neg echo intrvl: 0
Storage type : 2
Sess down time : 00:00:00
Sess Down Reason : NA
Bfd GTSM Disabled
Bfd Authentication Disabled

Counters values:
Pkt In : n/a                               Pkt Out : n/a
Pkts Drop : 00000000000000000000           Auth Pkts Drop : 00000000000000000000
Echo Out : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000001231         IPv6 Pkt Out : 00000000000000001233
UP Count : 1                               UPTIME : 00:04:30

Protocol Client Info:
ISIS-> Client ID: 6      Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

```
R1#
```

Originate Default Route to ISISv6 Neighbors

This example shows how to originate default route present to ISISv6 neighbors.

Note: To get a default route in ISIS, we must have it (said default route) in the routing table first.To get a default route in ISIS, we should have it (said default route) in the routing table first otherwise we can use the CLI "default information originate always" which advertises default route.

Topology

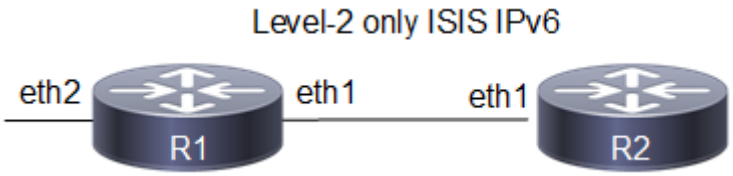

Figure 2-25: Basic IS-ISv6 Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#ipv6 route ::/0 2345::2`	Configure default route.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 address 2345::1/64`	Configure IPv6 address on interface.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).
`(config-router)#is-type level-2-only`	Configure instance as level-2-only routing.
`(config-router)#net 49.0001.0000.0000.0001.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#address-family ipv6`	Enter ipv6 address family.
`(config-router-af)#default-information originate`	Originate reachability information to Default destination into LSP.
`(config-router-af)#end`	Exit all modes and enter privilege mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router isis ABC`	Enable IS-ISv6 routing on an interface for area 49(ABC).

(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#is-type level-2-only	Configure instance as level-2-only routing.
(config-router)#net 49.0005.0000.0000.0002.00	Set a Network Entity Title for this instance, specifying the area address and the system ID.
(config-router)#end	Exit current mode and enter privilege mode

Validation

```
R1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag ABC:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     6         L2   IS-IS
R2#

R2#show ipv6 route isis
IP Route Table for VRF "default"
i L2   ::/0 [115/10] via fe80::ba6a:97ff:fecb:3ec5, eth1, 00:09:12
R2#

R2#show ipv6 isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag ABC:  VRF : default
L2   ::/0 [10]
      via fe80::ba6a:97ff:fecb:3ec5, eth1

R2#
R2#show isis database verbose
Tag ABC:  VRF : default
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0001.00-00  0x00000006   0x5FA4        1033              0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  Metric:   10         IS 0000.0000.0001.01
  Metric:   0          IPv6 ::/0
0000.0000.0001.01-00  0x00000001   0x21B9        628               0/0/0
  Metric:   0          IS 0000.0000.0001.00
  Metric:   0          IS R2.00
R2.00-00            * 0x00000002   0xFBED        633               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
```

```
Hostname:     R2
Metric:   10          IS 0000.0000.0001.01
```

CHAPTER 3 IS-IS Graceful Restart Configuration

The Intermediate System to Intermediate System (IS-IS) routing protocol is a link state intra-domain routing protocol. Normally, when an IS-IS router is restarted, temporary disruption of routing occurs due to events in both the restarting router and the neighbors of the restarting router.

ISIS provides graceful restart, in which the adjacency and routes are maintained in the routing table for the grace period. In this way, the data flow is not affected, and there is no packet loss during the restart phase.

With ISIS GR, the ISIS router should be able to restart gracefully with non-stop forwarding during the recovery. And the Helper ISIS router should be able to help restarting router by maintaining the adjacency.

ISIS Grace Restart Functionality applies to:

- ISIS broadcast network
- ISIS point-to-point network
- IPv4 domain
- IPv6 domain

Topology

In this example, R1 is the L1/L2 router, and R2 is the L1/L2 restart-helper router.

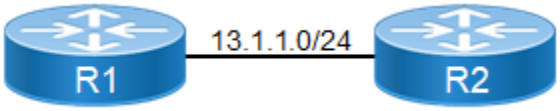

Figure 3-26: IS-IS Graceful Restart

Configuration

The following configuration is given only for R2, assuming that the adjacency with R1 is already up and the route tables with the appropriate routes are already populated.

R1

`#configure terminal`	Enter configure mode.
`(config)#isis restart grace-period 300`	Set the grace period to 300 seconds. The restarting router should come up before 300 seconds, otherwise, the adjacency and routes will be deleted.
`(config)#commit`	Commit candidate configuration to the running configuration

R2

`#configure terminal`	Enter configure mode.

`(config)#isis restart helper`	Configure this router as a restart helper.
`(config)#isis restart grace-period 100`	Set the grace period to 100 seconds. The restarting router should come up before 100 seconds, otherwise, the adjacency and routes will be deleted.

Note: The IS-IS daemon in the restarting router must be manually restarted using the restart IS-IS graceful command; it does not restart automatically.

Note: The unplanned GR is not supported.

Note: In the restarting node,

- restart the IS-IS daemon before it's LSP's remaining lifetime expires in Helper nodes.
- the restarting node might take higher graceful restart time interval than what it advertised (before GR) as 'Remaning Lifetime' in its LSP.
- in such cases, the LSP's received from the restarting node are not refreshed in helper nodes. These timed-out LSP's are deleted from Helper nodes and disrupt the traffic, though IS-IS adjacency is unaffected. To avoid this, the administrator has to ensure that GR time window is always lesser than the 'Remaning Lifetime' in Helper nodes.

Validation

R1

```
#restart isis graceful
% Warning :ISISD process will stop and needs to restart manually,
 You may loose ISIS configuration, if not saved
  Proceed for graceful restart? (y/n):y

Or
```

R1

```
#restart isis graceful grace-period 300
% Warning :ISISD process will stop and needs to restart manually,
 You may loose ISIS configuration, if not saved
  Proceed for graceful restart? (y/n):y
```

R2

```
R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     13.1.1.0/24 is directly connected, eth1, 04:08:20
i ia  20.0.0.0/6 [115/11] via 13.1.1.2, eth1, 00:10:44
```

```
i L1  33.0.0.0/24 [115/20] via 13.1.1.2, eth1, 00:10:44
C     127.0.0.0/8 is directly connected, lo, 04:10:59
C     192.168.52.0/24 is directly connected, eth0, 04:10:55


R2#show clns neighbors

Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 xe15        00e0.4bbe.ca02      Up     300       L2   IS-IS


R2#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  13.1.1.0/24 is directly connected, eth1, 04:10:56
i L1     13.1.1.0/24 [115/10] is directly connected, eth1, 01:58:50
i ia *>  20.0.0.0/6 [115/11] via 13.1.1.2, eth1, 00:13:20
i L1 *>  33.0.0.0/24 [115/20] via 13.1.1.2, eth1, 00:13:20
C    *>  127.0.0.0/8 is directly connected, lo, 04:13:35
C    *>  192.168.52.0/24 is directly connected, eth0, 04:13:31

Gateway of last resort is not set
```

CHAPTER 4 ISIS Multi Topology

Overview

Intermediate System to Intermediate System (ISIS) is a link-state routing protocol commonly used in large-scale service provider networks and enterprise networks. By default, ISIS is in a single topology with no separate Shortest Path First (SPF) process to differentiate between IPv4 and IPv6 topologies. If the topology in IPv6 is different from IPv4, the routing calculation encounters a problem as the routes are evaluated and chosen based on the common topology.

Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain. This means, both IPv4 and IPv6 have different topologies in the network and two SPF processes are run to find the route to each IPv4 and IPv6 destination independently.

Feature Characteristics

The main characteristics of ISIS Multi Topology are as follows:

- Enables ISIS to maintain separate topologies for IPv4 and IPv6 within the same ISIS area or domain.
- Allows routers in the ISIS area (for Level 1 routing) or domain (for Level 2 routing) to support both IPv4 and IPv6 address families.
- Performs multiple SPF calculations for each configured topology.
- Defines new Type-Length-Value (TLV) encodings called Multi Topology TLV (MT TLV). It is used to advertise the multiple topologies supported by the routers and contains information about the topology, including the ID (MTID), flags, and MT metric.
    - MT TLV (229): Capability TLV advertised in Hello packets.
    - MT intermediate system TLV (222): Extended TLV that describes the adjacency between nodes once the adjacency is formed.
    - MT IPV6 reachability TLV (237): Reachability TLV that gives information on IPv6 routing.

Benefits

The key benefits of ISIS Multi Topology are as follows:

- Enables the ability to make changes to the IPv6 topology without affecting the IPv4 topology, and vice-versa.
- Leverages common adjacency and database tables.
- Provides an independent SPF process for IPv4 and IPv6.

Prerequisites

- To enable ISIS Multi Topology on OcNOS devices, wide metric configuration is mandatory.
- Follow the below configuration steps to prepare the interface for implementation of Multi Topology by enabling single topology on the routers:

    Note: In each of the commands, modify the relevant router as R1, R2, R3, R4 or R5, depending on the router being configured.

1. Enter configure mode followed by interface mode on loopback interface.
```
#configure terminal
R1(config)#int lo
```
2. Configure the IP address for the interface.
```
R1(config -if)# ip add 1.1.1.1/32 secondary
R1(config -if)# ipv6 address 1111::11/128
```
3. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R1(config -if)# ip router isis 1
R1(config -if)# ipv6 router isis 1
R1(config -if)# exit
```
4. Enter the interface configuration mode and configure the IP address for the interface.
```
R1(config)#int xe22
R1(config -if)# ip address 10.1.1.1/24
R1(config -if)# ipv6 address 1001::1/64
```
5. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R1(config -if)# ip router isis 1
R1(config -if)# ipv6 router isis 1
R1(config -if)# exit
```

For Routers R1 and R5, continue the configuration steps as follows:

6. Set the routing process ID as 1 and configure the IS type as level-1.
```
R1(config)# router isis 1
R1(config-router)# is-type level-1
```
7. Configure wide metric-style.
```
R1(config-router)# metric-style wide
```
8. Enable dynamic host name under ISIS process.
```
R1(config-router)# dynamic-hostname
```
9. Enable BFD in all the interfaces.
```
R1(config-router)# bfd all-interfaces
```
10. Configure Network Entity Title (NET).
```
R1(config-router)# net 49.0000.0000.0001.00
```
11. Commit the candidate configuration to the running configuration.
```
R1(config-router)# commit
```

For Routers R2 and R4, use the following configuration steps after you exit the interface mode (step 5 shown above):

1. Enter the interface configuration mode and configure the IP address for the interface.
```
R2(config)#int xe24
R2(config -if)# ip address 20.1.1.1/24
R2(config -if)# ipv6 address 2001::1/64
```
2. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R2(config -if)# ipv6 router isis 1
R2(config -if)# exit
```
3. Enter the interface configuration mode and configure the IP address for the interface.
```
R2(config)#int xe23
```

```
R2(config -if)# ip address 40.1.1.1/24
R2(config -if)# ipv6 address 4001::1/64
```

4. Include the interface in the router's ISIS 1 instance and exit the interface mode.

```
R2(config -if)# ip router isis 1
R2(config -if)# ipv6 router isis 1
R2(config -if)# exit
```

5. Set the routing process ID as 1 and configure IS type as level 1.

```
R2(config)# router isis 1
R2(config-router)# is-type level-1
```

6. Configure wide metric style.

```
R2(config-router)# metric-style wide
```

7. Enable dynamic host name under ISIS process.

```
R2(config-router)# dynamic-hostname
```

8. Enable BFD in all the interfaces.

```
R2(config-router)# bfd all-interfaces
```

9. Configure Network Entity Title (NET).

```
R2(config-router)# net 49.0000.0000.0002.00
```

10. Commit the candidate configuration to the running configuration.

```
R2(config-router)# commit
```

For Router R3, follow these configuration steps after you exit the interface mode (step 5 shown above):

1. Enter Interface configuration mode and configure the IP address of the interface.

```
R3(config)#int xe31/1
R3(config -if)# ip address 50.1.1.1/24
R3(config -if)# ipv6 address 5001::1/64
```

2. Include the interface in the router's ISIS 1 instance and exit the interface mode.

```
R3(config -if)# ip router isis 1
R3(config -if)# ipv6 router isis 1
R3(config -if)# exit
```

3. Set the routing process ID as 1 and configure IS type as level-1.

```
R3(config)# router isis 1
R3(config-router)# is-type level-1
```

4. Configure wide metric-style.

```
R3(config-router)# metric-style wide
```

5. Enable dynamic host name under ISIS process.

```
R3(config-router)# dynamic-hostname
```

6. Enable BFD on all the interfaces.

```
R3(config-router)# bfd all-interfaces
```

7. Configure Network Entity Title (NET).

```
R3(config-router)# net 49.0000.0000.0003.00
```

8. Commit the candidate configuration to the running configuration.

```
R3(config-router)# commit
```

Configuration

To set up Multi Topology in ISIS, the configuration is as shown below:

Topology

This topology diagram consists of five routers (R1, R2,R3,R4 and R5).

It has both ISIS IPv4 and IPv6 routing enabled, except the link between R2 and R4 which has only IPv6 enabled.

In Single Topology, router R1 receives the information and calculates a SPF tree. To reach 5.5.5.5 (R5 IPv4), it takes the path R1-> R2 -> R4 ->R5. However, it fails since R2 to R4 is solely an IPv6 path. Since the same SPF tree is used for both IPv4 and IPv6 in R1, it considers the link between R2 -> R4 as the shortest path instead of R2 -> R3 -> R4.

On enabling Multi Topology on all the routers, SPF trees are calculated separately for IPv4 and IPv6 routing. This means, to reach from R1 to R5, IPv4 takes the R1 -> R2 -> R3 -> R4 -> R5 path and IPv6 takes the R1 -> R2 -> R4 -> R5 path.

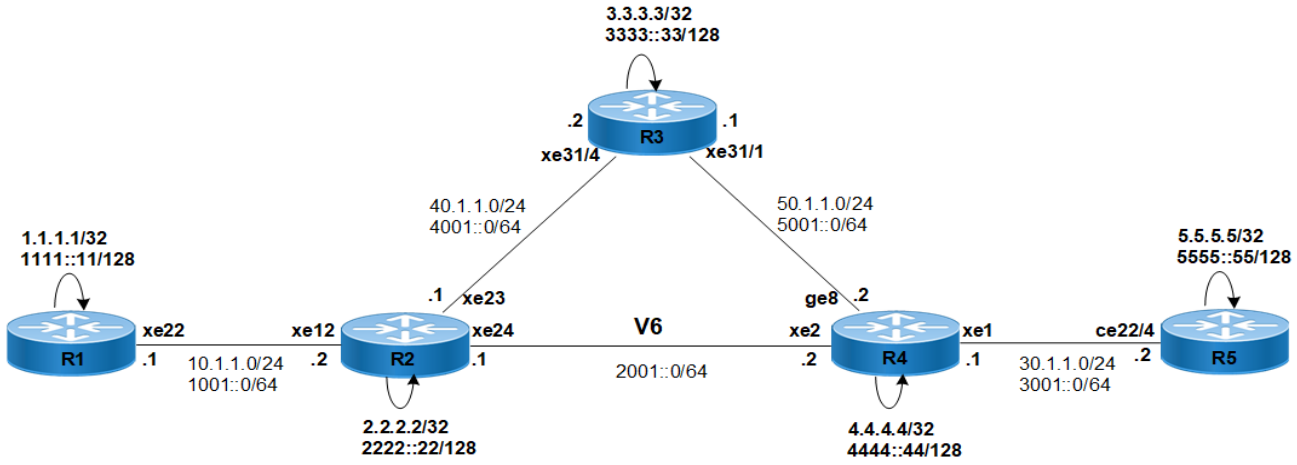

Figure 4-27: ISIS Multi Topology

To configure multi topology on the routers R1, R2, R3, R4 and R5, follow the steps mentioned below:

Note: Ensure that the Prerequisites are met for all the routers.

Note: Modify the commands for the relevant routers being configured (R1, R2, R3, R4 or R5).

1. Set the routing process ID as 1.

```
R1(config)# router isis 1
```

2. Configure metric-style wide.

```
R1(config-router)# metric-style wide
```

3. Configure address family IPv6.

```
R1(config-router)#address-family ipv6
```

4. Enable multi topology with level 1.

```
R1(config-router-af)#multi-topology level-1
```

5. Commit the candidate configuration to the running configuration.

```
R1(config-router-af)#commit
```

Validation for Multi Topology

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R2             xe22        00e0.4b77.39fe      Up     19        L1   M-ISIS

R1#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R2             xe22        Up     L1   64        0000.0000.0001.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:09:39
  Area Address(es): 49
  IP Address(es): 10.1.1.2
  IPv6 Address(es): fe80::2e0:4bff:fe77:39fe
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R1#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       xe22        00e0.4b77.39fe
R3                   20         R2       xe22        00e0.4b77.39fe
R4                   30         R2       xe22        00e0.4b77.39fe
R5                   40         R2       xe22        00e0.4b77.39fe


R1#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
```

```
R2                  10        R2       xe22        00e0.4b77.39fe
R3                  20        R2       xe22        00e0.4b77.39fe
R4                  20        R2       xe22        00e0.4b77.39fe
R5                  30        R2       xe22        00e0.4b77.39fe


R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, installed 01:55:53, last update
01:55:53 ago
i L1         2.2.2.2/32 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         3.3.3.3/32 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         4.4.4.4/32 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1         5.5.5.5/32 [115/50] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
C            10.1.1.0/24 is directly connected, xe22, installed 01:55:53, last update
01:55:53 ago
i L1         30.1.1.0/24 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1         40.1.1.0/24 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         50.1.1.0/24 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
C            127.0.0.0/8 is directly connected, lo, installed 01:57:14, last update
01:57:14 ago

Gateway of last resort is not set

R1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 01:57:15, last update 01:57:15 ago
C       1001::/64 via ::, xe22, installed 01:32:33, last update 01:32:33 ago
C       1111::11/128 via ::, lo, installed 01:33:09, last update 01:33:09 ago
```

```
i L1    2001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    3001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    3333::33/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    4001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    4444::44/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    5001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    5555::55/128 [115/40] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
C       fe80::/64 via ::, xe25, installed 01:56:18, last update 01:56:18 ago

R1#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:09:57.608 ago


R1#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000015   0x9E64        602               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00            * 0x0000000C   0x724E        602               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        601               0/0/0
  Area Address: 49
```

```
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        601               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
```

```
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        601               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        601               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        601               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R1#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000015   0x9E64        596               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
```

```
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00            * 0x0000000C   0x724E        596               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        595               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
```

```
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        595               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        595               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        595               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R2:

R2#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             xe12        e8c5.7a69.446f  Up     6         L1   M-ISIS
```

```
R3              xe23        903c.b3c5.ae9b      Up     6          L1   M-ISIS
R4              xe24        9819.2ccf.ede3      Up     9          L1   M-ISIS

R2#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R1             xe12        Up     L1   64        0000.0000.0001.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IP Address(es): 10.1.1.1
  IPv6 Address(es): fe80::eac5:7aff:fe69:446f
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R3             xe23        Up     L1   64        0000.0000.0003.03
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IP Address(es): 40.1.1.2
  IPv6 Address(es): fe80::923c:b3ff:fec5:ae9b
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R4             xe24        Up     L1   64        0000.0000.0004.04
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede3
  Topology: IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R2#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id           Metric     Next-Hop             Interface   SNPA
R1                  10         R1       xe12        e8c5.7a69.446f
R2                  --
```

```
R3                    10         R3       xe23        903c.b3c5.ae9b
R4                    20         R3       xe23        903c.b3c5.ae9b
R5                    30         R3       xe23        903c.b3c5.ae9b


R2#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id             Metric     Next-Hop             Interface   SNPA
R1                    10         R1       xe12        e8c5.7a69.446f
R2                    --
R3                    10         R3       xe23        903c.b3c5.ae9b
R4                    10         R4       xe24        9819.2ccf.ede3
R5                    20         R4       xe24        9819.2ccf.ede3

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/20] via 10.1.1.1, xe12, installed 01:11:03, last update
01:11:03 ago
C            2.2.2.2/32 is directly connected, lo, installed 01:59:20, last update
01:59:20 ago
i L1         3.3.3.3/32 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
i L1         4.4.4.4/32 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
i L1         5.5.5.5/32 [115/40] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C            10.1.1.0/24 is directly connected, xe12, installed 01:57:30, last update
01:57:30 ago
C            20.1.1.0/24 is directly connected, xe24, installed 01:59:19, last update
01:59:19 ago
i L1         30.1.1.0/24 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C            40.1.1.0/24 is directly connected, xe23, installed 01:59:19, last update
01:59:19 ago
i L1         50.1.1.0/24 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:20:04, last update
02:20:04 ago

Gateway of last resort is not set

R2#show ipv6 route
```

```
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:20:05, last update 02:20:05 ago
C       1001::/64 via ::, xe12, installed 01:32:42, last update 01:32:42 ago
i L1    1111::11/128 [115/20] via fe80::eac5:7aff:fe69:446f, xe12, installed 00:11:04,
last update 00:11:04 ago
C       2001::/64 via ::, xe24, installed 01:59:20, last update 01:59:20 ago
C       2222::22/128 via ::, lo, installed 01:33:21, last update 01:33:21 ago
i L1   3001::/64 [115/20] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04, last
update 00:11:04 ago
i L1    3333::33/128 [115/20] via fe80::923c:b3ff:fec5:ae9b, xe23, installed 01:11:04,
last update 01:11:04 ago
C       4001::/64 via ::, xe23, installed 01:24:52, last update 01:24:52 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
i L1   5001::/64 [115/20] via fe80::923c:b3ff:fec5:ae9b, xe23, installed 01:11:04, last
update 00:11:04 ago
                  [115/20] via fe80::9a19:2cff:fecf:ede3, xe24
i L1    5555::55/128 [115/30] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
C       fe80::/64 via ::, xe12, installed 01:57:31, last update 01:57:31 ago



R2#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:11:11.544 ago



R2#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
```

```
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        527               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00            * 0x00000014   0x2A52        528               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
```

```
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        527               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        527               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        527               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        527               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00
```

```
R2#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        520               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00            * 0x00000014   0x2A52        521               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
```

```
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        520               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        520               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        520               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        520               0/0/0
  Metric:   0          IS-Extended R5.00
```

```
  Metric:   0          IS-Extended R4.00


R3:



R3#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R4             xe31/1      9819.2ccf.ede9      Up     9         L1   M-ISIS
R2             xe31/4      00e0.4b77.3a09      Up     27        L1   M-ISIS


R3#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R4             xe31/1      Up     L1   64        0000.0000.0004.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:11:42
  Area Address(es): 49
  IP Address(es): 50.1.1.2
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede9
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
R2             xe31/4      Up     L1   64        0000.0000.0003.03
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:11:42
  Area Address(es): 49
  IP Address(es): 40.1.1.1
  IPv6 Address(es): fe80::2e0:4bff:fe77:3a09
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R3#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
```

```
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe31/4      00e0.4b77.3a09
R2                   10         R2       xe31/4      00e0.4b77.3a09
R3                   --
R4                   10         R4       xe31/1      9819.2ccf.ede9
R5                   20         R4       xe31/1      9819.2ccf.ede9


R3#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe31/4      00e0.4b77.3a09
R2                   10         R2       xe31/4      00e0.4b77.3a09
R3                   --
R4                   10         R4       xe31/1      9819.2ccf.ede9
R5                   20         R4       xe31/1      9819.2ccf.ede9

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/30] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1         2.2.2.2/32 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
C            3.3.3.3/32 is directly connected, lo, installed 02:00:27, last update
02:00:27 ago
i L1         4.4.4.4/32 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1         5.5.5.5/32 [115/30] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1        10.1.1.0/24 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1        30.1.1.0/24 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
C           40.1.1.0/24 is directly connected, xe31/4, installed 02:00:09, last update
02:00:09 ago
C           50.1.1.0/24 is directly connected, xe31/1, installed 02:00:26, last update
02:00:26 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:18:52, last update
02:18:52 ago

Gateway of last resort is not set
```

```
R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:18:53, last update 02:18:53 ago
i L1    1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    2001::/64 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
                  [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    3001::/64 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       3333::33/128 via ::, lo, installed 01:31:50, last update 01:31:50 ago
C       4001::/64 via ::, xe31/4, installed 01:30:10, last update 01:30:10 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       5001::/64 via ::, xe31/1, installed 01:29:43, last update 01:29:43 ago
i L1    5555::55/128 [115/30] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       fe80::/64 via ::, xe31/4, installed 02:00:10, last update 02:00:10 ago


R3#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:12:00.519 ago



R3#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
```

```
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        478               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00            * 0x00000013   0x7FCC        479               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
```

```
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00            * 0x0000000C   0x6D4E        479               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00             0x00000015   0x8C0D        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00             0x00000007   0x9A25        478               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00             0x0000000C   0x6751        478               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00             0x00000010   0xFA0F        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00             0x00000007   0xA813        478               0/0/0
  Metric:   0          IS-Extended R5.00
```

```
  Metric:   0          IS-Extended R4.00

R3#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        471               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00            * 0x00000013   0x7FCC        472               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
```

```
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00           * 0x0000000C   0x6D4E        472               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00             0x00000015   0x8C0D        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00             0x00000007   0x9A25        471               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00             0x0000000C   0x6751        471               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00             0x00000010   0xFA0F        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00             0x00000007   0xA813        471               0/0/0
```

```
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00



R4:

R4#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R5             xe1         e001.a6aa.0f23  Up     6         L1   M-ISIS
R2             xe2         00e0.4b77.3a0a  Up     22        L1   M-ISIS
R3             ge8         903c.b3c5.ae98  Up     22        L1   M-ISIS


R4#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R5             xe1         Up     L1   64        0000.0000.0005.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:38
  Area Address(es): 49
  IP Address(es): 30.1.1.2
  IPv6 Address(es): fe80::e201:a6ff:feaa:f23
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R2             xe2         Up     L1   64        0000.0000.0004.04
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:37
  Area Address(es): 49
  IPv6 Address(es): fe80::2e0:4bff:fe77:3a0a
  Topology: IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R3             ge8         Up     L1   64        0000.0000.0004.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:38
```

```
  Area Address(es): 49
  IP Address(es): 50.1.1.1
  IPv6 Address(es): fe80::923c:b3ff:fec5:ae98
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R4#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   30         R3       ge8         903c.b3c5.ae98
R2                   20         R3       ge8         903c.b3c5.ae98
R3                   10         R3       ge8         903c.b3c5.ae98
R4                   --
R5                   10         R5       xe1         e001.a6aa.0f23


R4#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe2         00e0.4b77.3a0a
R2                   10         R2       xe2         00e0.4b77.3a0a
R3                   10         R3       ge8         903c.b3c5.ae98
R4                   --
R5                   10         R5       xe1         e001.a6aa.0f23


R4#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/40] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
i L1         2.2.2.2/32 [115/30] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
i L1         3.3.3.3/32 [115/20] via 50.1.1.1, ge8, installed 01:01:13, last update
01:01:13 ago
C            4.4.4.4/32 is directly connected, lo, installed 02:01:55, last update
02:01:55 ago
```

```
i L1         5.5.5.5/32 [115/20] via 30.1.1.2, xe1, installed 01:12:47, last update
01:12:47 ago
i L1         10.1.1.0/24 [115/30] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
C            20.1.1.0/24 is directly connected, xe2, installed 02:01:04, last update
02:01:04 ago
C            30.1.1.0/24 is directly connected, xe1, installed 02:01:55, last update
02:01:55 ago
i L1         40.1.1.0/24 [115/20] via 50.1.1.1, ge8, installed 01:01:13, last update
01:01:13 ago
C            50.1.1.0/24 is directly connected, ge8, installed 02:01:22, last update
02:01:22 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:20:17, last update
02:20:17 ago

Gateway of last resort is not set


R4#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:20:18, last update 02:20:18 ago
i L1    1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48, last
update 00:12:48 ago
i L1    1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C       2001::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C       3001::/64 via ::, xe1, installed 01:33:20, last update 01:33:20 ago
i L1    3333::33/128 [115/20] via fe80::923c:b3ff:fec5:ae98, ge8, installed 01:01:14,
last update 01:01:14 ago
i L1    4001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 01:04:04, last
update 00:12:48 ago
                  [115/20] via fe80::923c:b3ff:fec5:ae98, ge8
C       4444::44/128 via ::, lo, installed 01:33:04, last update 01:33:04 ago
C       5001::/64 via ::, ge8, installed 01:29:27, last update 01:29:27 ago
i L1    5555::55/128 [115/20] via fe80::e201:a6ff:feaa:f23, xe1, installed 00:12:48,
last update 00:12:48 ago
C       fe80::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago


R4#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
```

```
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:12:55.361 ago



R4#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        423               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
```

```
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        423               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00            * 0x00000015   0x8C0D        424               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00            * 0x00000007   0x9A25        424               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00            * 0x0000000C   0x6751        424               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        423               0/0/0
```

```
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        423               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R4#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        417               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
```

```
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        417               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00            * 0x00000015   0x8C0D        418               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00            * 0x00000007   0x9A25        418               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
```

```
R4.04-00            * 0x0000000C   0x6751        418               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        417               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R5:

R5#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
R4             ce22/4      9819.2ccf.ede2      Up     28        L1   M-ISIS


R5#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R4             ce22/4      Up     L1   64        0000.0000.0005.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:13:32
  Area Address(es): 49
  IP Address(es): 30.1.1.1
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede2
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
```

```
R5#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   40         R4       ce22/4      9819.2ccf.ede2
R2                   30         R4       ce22/4      9819.2ccf.ede2
R3                   20         R4       ce22/4      9819.2ccf.ede2
R4                   10         R4       ce22/4      9819.2ccf.ede2
R5                   --


R5#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   30         R4       ce22/4      9819.2ccf.ede2
R2                   20         R4       ce22/4      9819.2ccf.ede2
R3                   20         R4       ce22/4      9819.2ccf.ede2
R4                   10         R4       ce22/4      9819.2ccf.ede2
R5                   --

R5#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/50] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1         2.2.2.2/32 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1         3.3.3.3/32 [115/30] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
i L1         4.4.4.4/32 [115/20] via 30.1.1.1, ce22/4, installed 01:13:40, last update
01:13:40 ago
C            5.5.5.5/32 is directly connected, lo, installed 02:03:15, last update
02:03:15 ago
i L1        10.1.1.0/24 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
C           30.1.1.0/24 is directly connected, ce22/4, installed 02:03:15, last update
02:03:15 ago
i L1        40.1.1.0/24 [115/30] via 30.1.1.1, ce22/4, installed 01:04:55, last update
01:04:55 ago
i L1        50.1.1.0/24 [115/20] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
```

```
C            127.0.0.0/8 is directly connected, lo, installed 02:20:59, last update
02:20:59 ago

Gateway of last resort is not set


R5#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:21:00, last update 02:21:00 ago
i L1    1001::/64 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    1111::11/128 [115/40] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    2001::/64 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    2222::22/128 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C       3001::/64 via ::, ce22/4, installed 01:05:32, last update 01:05:32 ago
i L1    3333::33/128 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    4001::/64 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    5001::/64 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C       5555::55/128 via ::, lo, installed 01:06:20, last update 01:06:20 ago
C       fe80::/64 via ::, ce22/4, installed 02:03:16, last update 02:03:16 ago


R5#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:13:45.938 ago



R5#show isis database verbose
Tag 1:  VRF : default
```

```
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        373               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        372               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
```

```
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        372               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        373               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        373               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00            * 0x00000010   0xFA0F        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
```

```
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00            * 0x00000007   0xA813        373               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00
```

Running Configuration

```
R1#sh running-config router isis
!
router isis 1
 is-type level-1
 metric-style wide
 dynamic-hostname
 bfd
 all-interfaces
 net 49.0000.0000.0001.00
 !
 address-family ipv6
 multi-topology
 level-1
 exit-address-family
!
R1#
```

CLI Commands

The ISIS Multi-topology feature introduces the `multi-topology` configuration command.

multi topology

Use this command to configure the ISIS topology type.

Use `no` parameter of this command to set the topology back to single.

Command Syntax

```
multi-topology (level-1|level-1-2|level-2)
no multi-topology
```

Parameters

`level-1`	Specify to enable multi-topology for level 1
`level-2`	Specify to enable multi-topology for level 2
`level-1-2`	Specify to enable multi-topology for both the levels

Default

ISIS topology type applies to levels 1 and 2.

Command Mode

Address-family IPv6 mode.

Applicability

Introduced the `multi-topology` parameter in OcNOS version 6.5.2.

Example

The following sequence of commands is used to configure ISIS `multi-topology` type as transition for levels 1 and 2.

```
(config)#router isis 1
(config-router)#address-family ipv6 unicast
(config-router-af)#multi-topology level-1-2
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
ISIS	Intermediate System to Intermediate System is a link-state routing protocol.
Multi Topology (MT)	In ISIS, Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain.
Type Length Value (TLV)	A data structure used to encode optional information in a data communications protocol: • Type: the kind of field that this part of the message represents • Length: the size of the value field, usually in bytes • Value: a variable-sized set of bytes that contains the data of the message
Shortest Path First (SPF)	Algorithm used by ISIS to make routing decisions based on the state of network links.
Loopback	A troubleshooting test in which a signal is transmitted from a source to a destination and then back to the source again so that the signal can be measured and evaluated.
Wide metric configuration	Allows ISIS to support larger networks by configuring high value metric in the interface.
Hello Packets	Information packets used to discover ISIS neighbors and maintain adjacencies.
Link State Packets (LSP)	Unidirectional, point-to-point, half-duplex connection used to exchange link state information.

Intermediate System to Intermediate System Command Reference

CHAPTER 1 IS-IS Commands

This chapter describes the IS-IS commands:

- accept-lifetime
- address-family ipv6
- adjacency-check
- area-password
- authentication key-chain
- authentication mode
- authentication send-only
- bfd all-interfaces
- capability cspf
- clear clns neighbors
- clear clns is-neighbors
- clear ip isis route
- clear isis adjacency
- clear isis counter
- clear isis interface counter
- clear isis process
- debug isis
- default-information originate
- distance (IPv4)
- distance (IPv6)
- domain password
- dynamic-hostname
- fast-reroute per-prefix
- fast-reroute terminate-hold-on interval
- fast-reroute tie-break
- ignore-lsp-errors
- ip router isis
- ip route high-priority tag
- ipv6 router isis
- isis authentication key-chain
- isis authentication mode md5
- isis authentication send-only
- isis bfd
- isis circuit-type
- isis csnp-interval

© 2025 IP Infusion Inc. Proprietary

- isis fast-reroute per-prefix candidate disable
- isis hello-interval
- isis hello-multiplier
- isis hello padding
- isis lsp-interval
- isis mesh-group
- isis metric
- isis network
- isis password
- isis priority
- isis retransmit-interval
- ispf
- isis wait-timer
- isis wide-metric
- isis tag
- isis te-metric
- is-type
- key chain
- key
- key-string
- key-string encrypted
- lsp-gen-interval
- lsp-mtu
- lsp-refresh-interval
- max-area-address
- max-lsp-lifetime
- metric-style
- net
- passive-interface
- prc-interval-exp
- redistribute
- redistribute isis
- redistribute isis WORD
- router isis
- send-lifetime
- set-overload-bit
- snmp restart isis
- spf-interval-exp
- summary-address

© 2025 IP Infusion Inc. Proprietary

- summary-prefix

accept-lifetime

Use this command to specify the time period during which the authentication on a key chain is received as valid.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> HH:MM:SS MONTH <01-31> <1993-2035>
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> infinite
accept-lifetime HH:MM:SS MONTH <01-31> <1993-2035> duration <1-2147483646>
no accept-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of accept-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to start. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to start as the first three letters of the month with first letter in upper case, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the end time of accept-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to end. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to end as the first three letters of the month with first letter in caps, for example, Jan. (case sensitive).
`<1993-2035>`	Specify the year to end.
`duration`	Indicate the duration parameter.
`<1-2147483646>`	Specify the actual end time duration of a key in seconds.
`infinite`	Specify the end time to never expire.

Default

By default, accept-lifetime command is disabled

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of accept-lifetime for key1 on the key chain named mychain.

```
#configure terminal
(config)#key chain mychain
```

```
(config-keychain)#key 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 03 2004 04:04:02 Oct 06
2006
```

address-family ipv6

Use this command to enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration.

Use the no parameter with this command to remove all configuration under 'address-family ipv6'.

Command Syntax

```
address-family ipv6 (unicast|)
no address-family ipv6 (unicast|)
```

Parameters

unicast	Specify unicast routing for IPv6.

Default

Unicast routing is not configured.

Command Mode

Router mode

Example

```
#configure terminal
(config)#router isis bb
(config-router)#address-family ipv6 unicast
```

adjacency-check

Use this command to check ISIS neighbor protocol support.

Use the `no` parameter with this command to uncheck ISIS neighbor protocol support.

Command Syntax

```
adjacency-check
no adjacency-check
```

Parameters

None

Default

By default, adjacency-check command is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis bb
(config-router)#adjacency-check

(config-router)#no adjacency-check
```

area-password

Use this command to set the authentication password for the Level-1 area and to set authentication on Level-1 sequence number PDUs. This command enables authentication when receiving and sending link-state and sequence number PDUs in Level-1 areas. The password must be the same for all the ISIS routers in the same area.

Use the `no` parameter with this command to clear the area password.

Command Syntax

```
area-password WORD
area-password WORD authenticate snp (send-only|validate)
no area-password
```

Parameters

`WORD`	Password string.
`authenticate`	Insert the password into Level-1 SNP PDUs.
`snp`	Sequence number PDUs.
`send-only`	Only insert the password into the Level-1 sequence number PDUs, but not check the password in sequence number PDUs that it receives. Use this keyword during a software upgrade to ease the transition.
`validate`	Insert the password into Level-1 sequence number PDUs and check the password in sequence number PDUs that it receives.

Default

By default, the area password is not configured

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis bb
(config-router)#area-password mypasswd

(config)#router isis bb
(config-router)#area-password myPass authenticate snp send-only

(config)#router isis bb
(config-router)#no area-password
```

authentication key-chain

Use this command to set the key chain to be used for authentication at the instance level. Authentication mode must be set to `md5` to configure the key chain. If no key chain is configured with the key-chain command, no key-chain authentication is performed.

Only one authentication key-chain is applied to an ISIS interface at a time. That is, issuing a second `isis authentication key-chain` command overrides the first `isis authentication key-chain` command. If neither the `level-1` nor the `level-2` keyword is configured, the chain applies to both levels. Authentication can be specified for an individual ISIS interface using the `isis authentication key-chain` command.

Use the `no` parameter with this command to unset the key chain used for authentication.

Command Syntax

```
authentication key-chain WORD (level-1|level-2|)
no authentication key-chain (level-1|level-2|)
```

Parameters

`WORD`	Specify the chain name (valid authentication keys).
`level-1`	Specify an authentication key-chain for level-1 PDUs.
`level-2`	Specify an authentication key-chain for level-2 PDUs.

Default

The key chain applies to the level(s) on which authentication mode is configured as MD5 if no level is specified.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router isis 1
(config-router)#authentication key-chain myKey level-1
```

authentication mode

Use this command to set the authentication mode at the instance level.

If clear-text authentication was configured using the `area-password` or `domain password` commands, the `authentication mode` command overrides both of those commands (based on the level at which MD5 is configured). If the `authentication mode` command was used first, and subsequently an attempt is made to use the `area-password` or `domain password` commands, the attempt fails. To configure clear-text authentication using the `area-password` or `domain password` commands, first use the `no authentication mode` command.

The type of authentication and the level to which it applies can be specified for a single ISIS interface, rather than per ISIS instance, using the `isis authentication mode` command.

Use the `no` parameter with this command to unset the authentication mode.

Command Syntax

```
authentication mode {md5|text} (level-1|level-2|)
no authentication mode {md5|text} (level-1|level-2|)
```

Parameters

`md5`	Keyed message digest
`text`	Text mode
`level-1`	Specify an authentication key-chain for level-1 PDUs.
`level-2`	Specify an authentication key-chain for level-2 PDUs.

Default

The authentication mode is set to MD5 for both levels if no level is specified.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#authentication mode md5 level-1

(config-router)#no authentication mode md5 level-1
```

authentication send-only

Use this command to set the send-only option at the instance level.

Use this command before configuring the authentication mode and authentication key-chain, so that the implementation of authentication goes smoothly. That is, the routers will have more time for the keys to be configured on each router if authentication is inserted only on the packets being sent, not checked on packets being received. After all routers that must communicate are configured with this command, enable the authentication mode and key chain on each router. Then, specify the `no authentication send-only` command to disable the send-only feature.

If the `level-1-2` keyword is configured, the send-only feature applies to both levels.

Use the `no` parameter with this command to unset the send-only option.

Command Syntax

```
authentication send-only (level-1-only|level-2-only|level-1-2)
no authentication send-only
```

Parameters

`level-1-only`	Set send-only option for level-1 only.
`level-2-only`	Set send-only option for level-2 only.
`level-1-2`	Set send-only option for level-1-2 only.

Default

By default, authentication send only is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#authentication send-only level-1-only

(config-router)#no authentication send-only
```

bfd all-interfaces

Use this command to enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance.

This command sets BFD fall-over check for all the neighbors under specified process. To disable BFD checking on particular interface use `isis bfd disable` command at interface mode.

Use the no parameter with this command to disable BFD functionality for an ISIS instance.

Command Syntax

```
bfd all-interfaces
no bfd all-interfaces
```

Parameters

None

Default

By default, the BFD feature is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis aa
(config-router)#bfd all-interfaces

(config-router)#no bfd all-interfaces
```

capability cspf

Use this command to enable the constrained shortest path first (CSPF) feature in the ISIS module. CSPF calculates optimum explicit route (ER), using Traffic Engineering Database and (TED) and pre-existing Label Switched Path (LSP). The resulting ER is used by a signaling protocol to set up LSPs. Use the no parameter with this command to disable CSPF functionality for an ISIS instance.

Command Syntax

```
capability cspf
no capability cspf
```

Parameters

None

Default

If this command is not used, the CSPF feature is disabled.

Command Mode

Router mode

Example

```
(config)#router isis aa
(config-router)#capability cspf
```

clear clns neighbors

Use this command to clear CLNS neighbor adjacencies.

Command Syntax

```
clear clns neighbors
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>ena
#clear clns neighbors
```

clear clns is-neighbors

Use this command to clear IS neighbor adjacencies.

Command Syntax

```
clear clns is-neighbors System-ID
```

Parameters

`System-ID`	Neighbor system ID in XXXX.XXXX.XXXX format.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#clear clns is-neighbors 1111.1111.1111.1111
```

clear ip isis route

Use this command to clear IPv4 routes.

Command Syntax

```
clear ip isis (WORD|) route (redistribution|all)
```

Parameters

`WORD`	Name that identifies the IS-IS area.
`redistribution`	Clear IS-IS local redistribution routes.
`all`	Clear all of the IS-IS routing tables.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>ena
#clear ip isis route redistribution
```

clear isis adjacency

This command is used to remove the entries from the IS-IS adjacency database. Clears all adjacencies for the specified routing instance or specified interface or particular system ID.

Command Syntax

```
clear isis adjacency (*|IFNAME|system-id XXXX.XXXX.XXXX|) (vrf VRFNAME|)
```

Parameters

`*`	Clear all neighbors.
`IFNAME`	Interface name.
`XXXX.XXXX.XXXX`	Neighbor System-ID.
`VRFNAME`	VRF name.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis adjacency *
```

clear isis counter

Use this command to clear system-wide IS-IS counters (IsisSystemCounterEntry in RFC 4444).

Command Syntax

```
clear isis counter
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis counter
```

clear isis interface counter

Use this command to clear interface counters. If you do not specify a parameter, then counters for all interfaces are cleared.

Command Syntax

```
clear isis interface counter (IFNAME|)
```

Parameters

IFNAME	Interface name.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis interface counter
```

clear isis process

Use this command to restart ISIS processes. If you do not specify a parameter, then all ISIS processes are restarted.

Command Syntax

```
clear isis (WORD|) process
```

Parameters

WORD	Name that identifies the IS-IS area.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis process
```

debug isis

Use this command to turn on debugging for specified criteria. Debug commands enable to show some debugging information about specified criteria into file or terminal.

Use the `no` parameter to turn off debugging for specified criteria.

Command syntax

```
debug isis (all|)
debug isis (authentication|bfd|checksum|events|hello|(interface IFNAME | System-ID|)|ifsm|local-updates|lsp|nfsm|nsm|pdu|protocol-errors|rib|spf)
no debug isis (all|)
no debug all
no debug all isis
no debug isis (authentication|bfd|checksum|events|hello|(interface IFNAME | System-ID|)|ifsm|local-updates|lsp|nfsm|nsm|pdu|protocol-errors|rib|spf)
```

Parameters

`all`	Enables all debugging.
`authentication`	Debugging for authentication.
`checksum`	Debugging for checksums.
`bfd`	Debugging for bidirectional forwarding detection.
`events`	Debugging for internal events.
`hello`	Debugging for hello processing.
`interface`	Interface.
`IFNAME`	Interface name.
`System-ID`	System identifier.
`ifsm`	Debugging for interface finite state machine.
`local-updates`	Debugging for local updates.
`lsp`	Debugging for link-state packet.
`nfsm`	Debugging for neighbor finite state machine.
`nsm`	Debugging for NSM messages.
`pdu`	Debugging for protocol data unit.
`protocol-errors`	Debugging for protocol errors.
`rib`	Debugging for RIB information.
`spf`	Debugging for shortest path first route calculation.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#debug isis pdu

#configure terminal
(config)#debug isis nsm
```

default-information originate

Use this command to originate reachability information to a default route into link-state packets.

There is no default information in a Level-2 domain by default, while a Level-1 router calculates a default to L1L2 route during SPF calculation. This command originates a default route into a Level-2 domain.

Use the `no` parameter with this command to withdraw reachability information to a default route from link-state packets.

Command Syntax

```
default-information originate (always| level-1|) (route-map WORD|)
no default-information originate (always| level-1|)
```

Parameters

`originate`	Specify to distribute a default route
`always`	The default route is advertised even if there is no default route in the router's routing table.
`level-1`	Distribute in level-1.
`route-map`	Identifies other filtering options via a route map.

Default

There is no default information in Level-2 domain by default, while Level-1 router calculates default to L1L2 route during SPF calculation. This command enables to originate default route into Level-2 domain. As an added option, if the user wants to originate the default route in L1 LSP, the "level-1" parameter can be used as follows:

```
default-information originate level-1
```

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#default-information originate
(config-router)#exit
(config)#exit
```

distance (IPv4)

Use this command in router mode to set the administrative distance for all IPv4 routes or routes from a specific source.

Use the no parameter with this command to remove an administrative distance.

Command Syntax

```
distance <1-255> (System-ID (WORD|)|)
no distance (System-ID|)
```

Parameters

`<1-255>`	Distance range.
`System-ID`	Source ID in XXXX.XXXX.XXXX format.
`WORD`	Access-list name.

Default

By default, all options are turned off.

Command Mode

Router mode

Examples

The following example shows setting the administrative distance for all routes.

```
#configure terminal
(config)#router isis
(config-router)#distance 10
```

The following example shows setting the administrative distance for a specific route source.

```
#configure terminal
(config)#router isis
(config-router)#distance 40 0000.0000.0001
```

distance (IPv6)

Use this command in router mode to set the administrative distance for all IPv6 routes.

Use the `no` parameter with this command to remove an administrative distance.

Command Syntax

```
distance <1-255>
no distance
```

Parameters

`<1-255>`	Distance range.

Default

By default, all options are turned off.

Command Mode

Address-family ipv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the administrative distance for all routes.

```
#configure terminal
(config)#router isis
(config-router)#address-family ipv6
(config-router-af)#distance 14
```

domain password

Use this command to set the authentication password for the Level-2 domain, and optionally, the authentication password on Level-2 sequence number PDUs.

Configuring this command to enable authentication when receiving and sending link state and sequence number PDUs in Level-2 domain. The domain password must be the same in the Level-2 domain.

Use the `no` parameter with this command to clear the domain password.

Command Syntax

```
domain password WORD
domain password WORD authenticate snp (send-only|validate)
no domain password
```

Parameters

`WORD`	Password string.
`authenticate`	Inserts th password into Level-1 sequence number PDUs.
`snp`	sequence number PDUs.
`send-only`	Only inserts the password into the Level-1 sequence number PDUs, but does not check the password in sequence number PDUs that it receives. Use this keyword during a software upgrade to ease the transition.
`validate`	Inserts the password into the Level-1 sequence number PDUs and checks the password in sequence number PDUs received.

Default

By default, there is no domain password.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#domain password mypasswd

(config)#router isis bb
(config-router)#domain password myPass authenticate snp send-only

(config)#router isis bb
(config-router)#no domain password
```

dynamic-hostname

Use this command to configure the hostname to advertise for an ISIS instance using the dynamic hostname exchange mechanism (RFC 2763) and system-ID-to-hostname translation. This command configures a hostname to use for the Dynamic Hostname Exchange Mechanism and System-ID to hostname translation. This is required to get accurate results when using the show isis database and a few other commands.

Use the `no` parameter to disable the Hostname configured.

Command Syntax

```
dynamic-hostname
hostname dynamic
dynamic-hostname area-tag
no dynamic-hostname
no hostname dynamic
```

Parameters

`area-tag`	Use the routing area tag as the hostname, not the router's global hostname.

Default

By default, the Dynamic Hostname Exchange Mechanism is disabled.

Command Mode

Router mode

Example

```
#configure terminal
(config)#router isis bb
(config-router)#dynamic-hostname area-tag
```

fast-reroute per-prefix

Use this command to enable Loop Free Alternate Fast Reroute (LFA FRR) for all prefixes or only those prefixes in a route map.

Use the no form of this command to disable LFA FRR.

Command Syntax

```
fast-reroute per-prefix (level-1|level-2) proto (ipv4) (all|route-map WORD)
 no fast-reroute per-prefix (level-1|level-2) proto (ipv4)
```

Parameters

`level-1`	Level 1 only.
`level-2`	Level 2 only.
`ipv4`	IPv4 address family only.
`all`	All prefixes.
`route-map`	Prefixes from a route map.
`WORD`	Route map name.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#fast-reroute per-prefix level-2 proto ipv4 all
(config-router)#fast-reroute per-prefix level-1 proto ipv4 route-map rmap1
(config-router)#no fast-reroute per-prefix level-2 proto ipv4
```

fast-reroute terminate-hold-on interval

Use this command to set the Loop Free Alternate Fast Reroute (LFA FRR) termination hold-on timer.

Use the no form of this command to set the termination hold-on timer to its default value (1000 milliseconds).

Command Syntax

```
fast-reroute terminate-hold-on interval <100-100000>
no fast-reroute terminate-hold-on interval
```

Parameters

<100-100000>	LFA FRR termination hold-on timer interval in milliseconds.

Default

1000 milliseconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#fast-reroute terminate-hold-on interval 7000
(config-router)#no fast-reroute terminate-hold-on interval
```

fast-reroute tie-break

Use this command to set the tie-breaking policy for selecting a fast reroute repair path. You assign a priority to each type of repair path. The tie-breaker value is used to select an LFA FRR route when multiple LFA FRR routes are available for the same primary route.

Use the `no` form of this command to set the tie-break preference value for a protection type to its default value.

To set all types of repair paths to their default priorities, do not specify a repair path with the `no` form of this command.

Command Syntax

```
fast-reroute tie-break (level-1|level-2) proto (ipv4) (primary-path|interface-
  disjoint|node-protecting|broadcast-interface-disjoint|secondary-path|downstream-
  path) index <1-255>
no fast-reroute tie-break (level-1|level-2) proto (ipv4) (primary- path|interface-
  disjoint|node-protecting|broadcast-interface-disjoint|secondary-
  path|downstream-path)
```

Parameters

`level-1`	Level 1 only.
`level-2`	Level 2 only.
`ipv4`	IPv4 address family only.
`primary-path`	Use a path from the Equal-Cost Multipath Path (ECMP) set. An ECMP found during the primary shortest path first (SPF) repair might not be desirable in networks where traffic exceeds the capacity of any single link.
`interface-disjoint`	Link protecting: prefer a backup path that uses a different interface than the interface used to reach destination via the primary path.
`node-protecting`	Bypass the `primary-path` gateway router which might not protect the router that is the next hop in the primary path. This ensures complete traffic protection even if the primary next-hop router fails.
`broadcast-interface-disjoint`	Do not use the interface if connected to a broadcast network. Repair paths protect links when a repair path and a protected primary path use *different* next-hop interfaces. However, on broadcast interfaces, if the repair path is computed via the same interface as the primary path, but their next-hop gateways are different, the router is protected but the link might not be.
`secondary-path`	Prefer a non-ECMP backup path.
`downstream-path`	Prefer a backup path to the destination which satisfies the downstream condition where the path cost to reach the destination from the LFA next hop is less than the path cost to the destination from the self node via primary next hop: Distance_opt(N, D) < Distance_opt(S, D) This might result in lost traffic, but prevents looping.
`index`	Tie breaking index. A lower value has higher preference.

`<1-255>` Tie breaking index value.

Defaults

`primary-path`	20
`interface-disjoint`	60
`node-protecting`	30
`broadcast-interface-disjoint`	70
`secondary-path`	255
`downstream-path`	90

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#fast-reroute tie-break level-2 proto ipv4 node-protecting
index 127
(config-router)#no fast-reroute tie-break level-1 proto ipv4
broadcastinterface-disjoint
```

ignore-lsp-errors

Use this command to ignore link-state packets (LSPs) with checksum errors. By default, ISIS validates the checksum for LSP and if the checksum has an error, the LSP is dropped. Giving this command says to ignore the LSP checksum error and treat it as if checksum is passed.

Use the `no` parameter to turn off this function.

Command Syntax

```
ignore-lsp-errors
no ignore-lsp-errors
```

Parameters

None

Default

By default, the LSP checksum is checked on receipt.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

In this sample, rtr1 does not drop LSP packets with bad checksum.

```
#configure terminal
(config)#router isis bb
(config-router)#ignore-lsp-errors
```

ip router isis

Use this command to enable ISIS IPv4 routing on the interface. This command is mandatory to ISIS configuration.

After giving this command, the router sends ISIS Hello with IP address TLV on this interface and IP reachability information TLV in link-state packets are updated.

Use the `no` parameter with this command to disable ISIS IPv4 routing on the interface. This action does not clear the ISIS database. To clear the database, unconfigure the ISIS routing instance.

Command Syntax

```
ip router isis WORD
no ip router isis WORD
```

Parameters

`WORD`	Name that identifies the IS-IS area. Specify an existing area name or a new area name.

Default

By default, IPv4 routing is disabled on the router.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ip router isis bb
```

ip route high-priority tag

Use this command to set the tag value on ISIS LSP.

Command Syntax

```
ip route high-priority tag <1-4294967295>
```

Parameters

`tag <1-4294967295>` Specify the high priority tag value.

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS Version 6.0.0.

Example

```
OcNOS#configure terminal
OcNOS(config)#router isis 1
OcNOS(config-router)#ip route high-priority tag 10
```

ipv6 router isis

Use this command to enable ISIS IPv6 routing on the interface. This command is mandatory to IPv6 ISIS configuration. Match the ISIS instance tag to one of existing instance's tags, or a new instance with the tag name should be initiated, otherwise routing will not run on this interface.

Configuring this command, the router sends ISIS Hello with IPv6 address TLV on this interface, and IPv6 reachability information TLV in the LSP will be updated.

Use the no parameter with this command to disable ISIS IPv6 routing on the interface.

Command Syntax

```
ipv6 router isis WORD
no ipv6 router isis WORD
```

Parameters

WORD	ISIS instance name.

Default

By default, IPv6 routing is disabled on the router.

Command Mode

Interface mode

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 router isis bb
```

isis authentication key-chain

Use this command to set the key chain to be used for authentication on the interface-related packets.

Authentication mode must be set to `md5` to configure the key chain. If no key chain is configured with the `key-chain` command, no key-chain authentication is performed. Only one authentication key-chain is applied to an ISIS interface at a time. That is, issuing a second `isis authentication key-chain` command overrides the first `isis authentication key-chain` command.

If neither the `level-1` nor `level-2` keyword is configured, the key chain applies to the level(s) on which the authentication mode is configured as `md5`. Authentication can be specified for an entire instance of ISIS, instead of at the interface level, by using the `authentication key-chain` command.

Use the `no` parameter with this command to unset the key chain used for authentication on the interface-related packets.

Command Syntax

```
isis authentication key-chain WORD (level-1|level-2|)
no isis authentication key-chain (level-1|level-2|)
```

Parameters

`WORD`	Chain name - valid authentication keys.
`level-1`	Specify an authentication key-chain for level-1 PDUs.
`level-2`	Specify an authentication key-chain for level-2 PDUs.

Default

By default, this option is disabled. The key chain applies to the level(s) on which authentication mode is configured as MD5 if no level is specified.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#isis authentication key-chain myKey level-1
```

isis authentication mode md5

Use this command to set the MD5 authentication mode. If the `isis password` command was used, then subsequently an attempt is made to use the `isis authentication mode` command, the attempt fails. To configure `isis authentication mode` command, first use the `no isis password` command. If the `isis authentication mode` command was used, then subsequently an attempt is made to use the `isis password` command, the attempt fails.

To configure clear text authentication using the `isis password` command, first use the `no isis authentication mode` command. The type of authentication and the level to which it applies can be specified for the entire ISIS instance, rather than per interface, using the `authentication mode` command.

Use the `no` parameter with this command to unset the MD5 authentication mode.

Command Syntax

```
isis authentication mode {md5|text} (level-1|level-2|)
no isis authentication mode {md5|text} (level-1|level-2|)
```

Parameters

`md5`	Keyed message digest
`text`	Text mode
`level-1`	Specify an authentication key-chain for level-1 PDUs.
`level-2`	Specify an authentication key-chain for level-2 PDUs.

Default

By default, this option is disabled. The authentication mode will be set to MD5 for both levels if no level is specified.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#isis authentication mode md5 level-1
```

isis authentication send-only

Use this command to set the send-only option to the interface-related packets.

Use this command before configuring the ISIS authentication mode and ISIS authentication key-chain, so that the implementation of authentication goes smoothly. That is, the routers will have more time for the keys to be configured on each router if authentication is inserted only on the packets being sent, not checked on packets being received. After all routers that must communicate are configured with this command, enable the authentication mode and key chain on each router.

Use the `no` parameter with this command to unset the send-only option to the interface-related packets.

Command Syntax

```
isis authentication send-only (level-1-only|level-2-only|level-1-2)
no isis authentication send-only
```

Parameters

`level-1-only`	Set send-only option for level-1 only.
`level-2-only`	Set send-only option for level-2 only.
`level-1-2`	Set send-only option for level-1-2 only.

Default

By default, this option is disabled. The send-only option applies to both levels if no level is specified.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#isis authentication send-only level-1-only
```

isis bfd

Use this command to enable/disable the BFD check on interface.

The `isis bfd` command allows a user to enable BFD on an interface. The `isis bfd disable` command disables BFD checking on an interface. However, the `no isis bfd` and `no isis bfd disable` commands both remove the enable/disable configuration, but do not disable/enable BFD.

The `bfd all-interfaces` command enables BFD on all interfaces attached to an instance then configuring. This command disables BFD configuration on a particular interface.

Command Syntax

```
isis bfd (disable|)
no isis bfd (disable|)
```

Parameters

`disable`	Specify to disable BFD.

Default

By default, bfd feature enable/disable is not configured.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis bfd disable
```

isis circuit-type

Use this command to set the circuit type for the interface.

If level-1 or level-2-only is specified in this command, ISIS sends only the specified level of PDUs. On the point-to-point interface, there is only one type of Hello packet, so in this case ISIS Hello will be sent regardless of circuit-type. If is-type is configured as level-1 or level-2 only, routing for this instance is performed for only the specified level. In this manner, only the particular level of PDU is sent on the interface.

Use the `no` parameter to reset circuit type to the default.

Command Syntax

```
isis circuit-type (level-1|level-1-2|level-2-only)
no isis circuit-type
```

Parameters

`level-1`	Specify that only Level-1 adjacencies are formed.
`level-1-2`	Specify that Level-1-2 adjacencies are formed.
`level-2-only`	Specify that only Level-2 adjacencies are formed.

Default

By default, the default circuit-type is level-1-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis circuit-type level-2-only
```

isis csnp-interval

Use this command to set CSNP (Complete sequence number PDU) interval in seconds.

Configuring this command changes the interval between two consecutive CSNP transmission. By default, CSNP is sent every 10 seconds only by LAN DIS. This parameter is only valid on broadcast interface, since periodic CSNP is only sent on broadcast interface, while CSNP on Point-to-Point interface is sent only when adjacency is initiated.

Use the `no` parameter with this command to reset CSNP interval to the default value.

Command Syntax

```
isis csnp-interval <1-65535> (level-1|level-2|)
no isis csnp-interval (level-1|level-2|)
```

Parameters

`<1-65535>`	Specify the CSNP interval in seconds.
`level-1`	Specify Level-1 CSNP.
`level-2`	Specify Level-2 CSNP.

Default

By default, ISIS uses 10 seconds for the interval and the interval is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis csnp-interval 20
```

isis fast-reroute per-prefix candidate disable

Use this command to prevent an interface from becoming an Loop Free Alternate Fast Reroute (LFA FRR) for a primary route.

Use the no form of this command to enable an interface to become an LFA FRR for a primary route.

Command Syntax

```
isis fast-reroute per-prefix candidate disable (level-1|level-2)
no isis fast-reroute per-prefix candidate disable (level-1|level-2)
```

Parameters

`level-1`	Level-1 only
`level-2`	Level-2 only

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis fast-reroute per-prefix candidate disable level-2
(config-if)#no isis fast-reroute per-prefix candidate disable level-2
```

isis hello-interval

Use this command to set the Hello interval in seconds. The Hello-interval is set with the hello-multiplier (see `isis hello-multiplier` command).

Configuring this command changes the time interval between two consecutive Hello transmissions. If a device receives its own LSP with a maximum sequence number, then it suspends ISIS for the hold interval. DIS sends Hello transmissions at three times the rate than non-DIS. If ISIS is elected as DIS on this interface, ISIS sends Hello every 3.3 seconds.

If `minimal` keyword is specified, Holding timer in Hello PDU is set to 1 second and Hello interval is calculated by dividing by the hello-multiplier. For example, if the hello-multiplier is configured as 4 and `hello-interval minimal` is the command used, an Hello PDU is sent every 250 milliseconds.

Use the `no` parameter to set the Hello interval to the default.

Command Syntax

```
isis hello-interval <1-65535> (level-1|level-2|)
isis hello-interval minimal (level-1-only|level-2-only|level-1-2)
no isis hello-interval (level-1|level-2|)
no isis hello-interval minimal
```

Parameters

`<1-65535>`	Specify the hello interval in seconds.
`minimal`	Specify the holding-time as 1 second.
`level-1`	Specify Level-1 CSNP.
`level-2`	Specify Level-2 CSNP.
`level-1-only`	Specify only Level-1 CSNP.
`level-2-only`	Specify only Level-2 CSNP.
`level-1-2`	Specify only Level-1-2 CSNP.

Default

By default, ISIS uses 10 seconds for the interval and the interval is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#isis hello-interval 5 level-1

(config-if)#isis hello-interval minimal level-1-only
```

isis hello-multiplier

Use this command to set multiplier for Hello holding time.

Changes Holding Timer in Hello PDU. Holding timer is calculated by “Hello-Interval” multiplied by this value. If `minimal` keyword is specified with the Hello-Interval, the holding timer is set to 1 second and the hello-interval is calculated by dividing 1 by this value.

Use the `no` parameter with this command to set multiplier to the default.

Command Syntax

```
isis hello-multiplier <2-100> (level-1|level-2|)
no isis hello-multiplier (level-1|level-2|)
```

Parameters

`<2-100>`	Specify a hello multiplier value.
`level-1`	Specify Level-1 hello.
`level-2`	Specify Level-2 hello.

Default

By default, ISIS uses 3 seconds for the multiplier value and the multiplier is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis hello-multiplier 4
```

isis hello padding

Use this command to enable IS-IS hello padding at the interface level.

Intermediate System-to-Intermediate System (IS-IS) hellos are padded to the full maximum transmission unit (MTU) size. The benefit of padding IS-IS hellos to the full MTU is that it allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.

Use the `no` parameter with this command to disable IS-IS hello padding.

Command Syntax

```
isis hello padding
```

```
no isis hello padding
```

Parameters

None

Default

By default, ISIS hello padding is enabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth1
(config-if)#isis hello padding

(config-if)#no isis hello padding
```

isis lsp-interval

Use this command to set the link-state packet (LSP) transmission interval.

Giving this command changes the minimum interval between two consecutive LSP transmission. When flooding or some other event triggers LSP transmission, the LSP is put in the interface queue and scheduled to send according to this interval. Two consecutive LSP transmissions are scheduled to have at least this interval.

Use the `no` parameter with this command to set LSP transmission interval to the default.

Command Syntax

```
isis lsp-interval <1-4294967295>
no isis lsp-interval
```

Parameters

`<1-4294967295>` LSP transmission interval in milliseconds.

Default

By default, ISIS uses 33 milliseconds for the interval.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#isis lsp-interval 100

(config-if)#no isis lsp-interval
```

isis mesh-group

Use this command to set Mesh Group ID on the current interface.

Use the `no` parameter to unset mesh group on the current interface.

Command Syntax

```
isis mesh-group <1-4294967295>
no isis mesh-group
```

Parameters

`<1-4294967295>` Specify a mesh group number

Default

By default, mesh groups are not enabled on this interface.

Command Mode

Interface mode

Examples

```
(config)#interface eth0
(config-if)#isis mesh-group 20

(config)#interface eth2
(config-if)#no isis mesh-group
```

isis metric

Use this command to set the default metric for the interface. The interface default metric is put into IP reachability information TLVs and in IS reachability information TLVs in link-state packets. The value is used for SPF calculation, and is applied when the metric-style is configured as “narrow”.

Use the `no` parameter with this command to set default metric to the default.

Command Syntax

```
isis metric <1-63> (level-1|level-2|)
no isis metric (level-1|level-2|)
```

Parameters

`<1-63>`	Default metric.
`level-1`	Default metric for level-1 circuit.
`level-2`	Default metric for level-2 circuit.

Default

By default, ISIS uses 10 for the metric value and the value is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis metric 20
```

isis network

Use this command to change a broadcast interface network type to a point-to-point network type.

Use the `no` parameter with this command to revert to the default setting of a broadcast interface network type.

Command Syntax

```
isis network (broadcast|point-to-point)
no isis network
```

Parameters

`broadcast`	Specify ISIS a broadcast multi-access network.
`point-to-point`	Specify ISIS a point-to-point network.

Default

By default, the network is set to a broadcast multi-access network type.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis network point-to-point
```

isis password

Use this command to set the authentication password of Hello PDU on the interface.

If the `isis authentication mode` command was used, then subsequently an attempt is made to use the `isis password` command, the attempt fails. To configure clear text authentication using the `isis password` command, first use the `no isis authentication mode` command.

Use the `no` parameter to clear the password.

Command Syntax

```
isis password WORD (level-1|level-2|)
no isis password (level-1|level-2|)
```

Parameters

`WORD`	Specify a password string.
`level-1`	Specify a password for Level-1 hello PDUs.
`level-2`	Specify a password for Level-2 hello PDUs.

Default

By default, no password is configured; this applies to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis password mypassword level-1
```

isis priority

Use this command to set the priority for LAN DIS election. This command changes the priority value in LAN ISIS Hello PDUs. A lower priority value is less preferred in DIS election, and a higher priority value is more preferred.

Note: This command is not valid for Point-to-Point interface.

Use the `no` parameter to set priority to the default.

Command Syntax

```
isis priority <0-127> (level-1|level-2|)
no isis priority (level-1|level-2|)
```

Parameters

`<0-127>`	Priority value
`level-1`	Specify a password for Level-1 hello PDUs.
`level-2`	Specify a password for Level-2 hello PDUs.

Default

By default, ISIS uses 64 for the priority value, and the priority is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#isis priority 127
```

isis retransmit-interval

Use this command to set the link-state packet (LSP) retransmission interval.

Use the `no` parameter to set the interval to the default.

Command Syntax

```
isis retransmit-interval <1-65535>
no isis retransmit-interval
```

Parameters

`<1-65535>`	Interval for retransmission of the same LSP in seconds.

Default

By default, ISIS uses an interval of 5 seconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#isis retransmit-interval 10

(config-if)#no isis retransmit-interval
```

ispf

Use this command to enable incremental SPF for a routing process.

Use the `no` parameter to disable incremental SPF from a routing process.

Command Syntax

```
ispf (level-1|level-2-only|)
no ispf
```

Parameters

`level-1`	Act as level-1 only IS.
`level-2-only`	Act as level-2 only IS.

Default

By default, all levels are turned off.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#ispf level-1

(config)#router isis bb
(config-router)#no ispf
```

isis wait-timer

Use the isis wait-timer seconds command to configure the number of seconds the router will wait for adjacency information

Use the `no` parameter to set the wait-timer to the default.

Command Syntax

```
isis wait-timer <1-65535> (level-1|level-2|)
no isis wait-timer (level-1|level-2|)
```

Parameters

`level-1`	Act as level-1 only IS.
`level-2`	Act as level-2 only IS.

Default

By default, wait-timer will be 20 seconds (2 times the hello timer).

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#isis wait-timer 30

(config)#router isis bb
(config-router)#no isis wait-timer
```

isis wide-metric

Use this command to set wide metric for the interface.

Interface wide-metric is put into Extended IP reachability TLVs. Extended IS reachability TLVs and IPv6 rechability TLVs in LSPs. The value is used for SPF calculation. The value is applied when metric-style is configured as 'wide'.

Use the `no` parameter to set wide metric to the default.

Command Syntax

```
isis wide-metric <1-16777214> (level-1|level-2|)
no isis wide-metric (level-1|level-2|)
```

Parameters

`<1-16777214>`	Specify a wide metric.
`level-1`	Specify the wide metric for Level-1 circuit.
`level-2`	Specify the wide metric for Level-2 circuit.

Default

By default, ISIS uses 10 for the metric value and the metric is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

None.

Examples

```
#configure terminal
(config)#interface eth0
(config-router)#isis wide-metric 100

(config)#interface eth0
(config-router)#no isis wide-metric
```

isis tag

Use this command to sets the tag for link-state packets (LSPs) sent out advertising routes for networks directly connected to an interface.

If you do not specify a parameter, then the tag value is set for level-1-2 boundary.

Use the `no` parameter to unset the tag.

Command Syntax

```
isis tag <1-4294967295> (level-1|level-2|)
no isis tag (level-1|level-2|)
```

Parameters

`<1-4294967295>`	Tag value.
`level-1`	Specify the tag value for the level-1 boundary.
`level-2`	Specify the tag value for the level-2 boundary.

Command Mode

Interface mode

Examples

```
>ena
#con term
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#interface eth0
(config-if)#isis tag 500 level-1
```

isis te-metric

Use this command to set Traffic Engineering (TE) metric for the interface. Interface te-metric is put into Extended IS reachability TLV. The value is used for CSPF calculation. The value is applied when metric-style is configured as wide.

Use the `no` parameter to unset the Traffic Engineering metric.

Command Syntax

```
isis te-metric <1-16777214> (level-1|level-2|)
no isis te-metric (level-1|level-2|)
```

Parameters

`<1-16777214>`	Specify a TE metric.
`level-1`	Specify the TE metric for level-1 circuit.
`level-2`	Specify the TE metric for level-2 circuit.

Default

By default, it will take ISIS wide metric value.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS-DC version 5.0

Examples

```
#configure terminal
(config)#interface eth1
(config-router)#isis te-metric 100
(config)#interface eth1
(config-router)#no isis te-metric
```

is-type

Use this command to set the IS to the specified level of routing.

Changing `is-type` brings down, then brings up a particular level of routing. There is a limitation: Only one ISIS instance can run Level-2 routing (either Level-2 only IS, or Level-1-2 IS).

Use the `no` parameter to set the IS to the default.

Command Syntax

```
is-type (level-1|level-1-2|level-2-only)
no is-type
```

Parameters

`level-1`	Act as level-1 only IS.
`level-1-2`	Act as level-1-2 IS.
`level-2-only`	Act as level-2 only IS.

Default

By default, ISIS uses level-1-2 if there is no Level-2 instance nor a Level-1-2 instance. Otherwise, it uses level-1.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#is-type level-1

(config)#router isis bb
(config-router)#no is-type
```

key chain

Use this command to enter the key chain management mode and to configure a key chain with a key chain name. This command allows you to enter the keychain mode to specify keys on this key chain.

Command Syntax

```
key chain WORD
no key chain WORD
```

Parameters

WORD	Specify the name of the key chain to manage.

Default

By default, keychain mode is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the creation of a key chain named `mychain` and the change to keychain mode:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#
```

key

Use this command to manage, add or delete authentication keys in a key-chain. This command allows you to enter the keychain-key mode to set a password for the key.

Command Syntax

```
key <0-2147483647>
no key <0-2147483647>
```

Parameters

`<0-2147483647>` Specify a key identifier.

Default

By default, ISIS uses level-1-2 if there is no Level-2 instance nor a Level-1-2 instance. Otherwise, it uses level-1.

Command Mode

Keychain mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example configures a key number 1 and shows the change to keychain-key command mode.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#
```

key-string

Use this command to define a password in plain-text to be used by a key.

The password is stored as encrypted, and is displayed in encrypted text when show running-config command is executed.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
key-string WORD
no key-string
```

Parameters

`WORD`	Specify a string of characters to be used as a password by the key. The length of the string should be between 1-80 characters.

Default

By default, password is not configured.

Command Mode

Keychain mode and Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#key-string prime

(config-keychain)#key 1
(config-keychain-key)#no key-string
```

key-string encrypted

Use this command to define a password in its encrypted format to be used by a key.

Use the `no` parameter with this command to disable this feature

Command Syntax

```
key-string encrypted WORD
no key-string encrypted
```

Parameters

`WORD`	Specify the encrypted string of characters to be used as a passoword by the key. The length of this string should be between 18-162 characters.

Default

By default, password is not configured.

Command Mode

Keychain mode and Keychain-key mode.

Applicability

This command was introduced in OcNOS version 4.1.

Examples:

In the following example, the encrypted password for key 1 in the key chain named mykeychain is set to 0xd6c50b442de47f70 (equivalent to "mychain" in plain-text):

```
#configure terminal
(config)#key chain mykeychain
(config-keychain)#key 1
(config-keychain-key)#key-string encrypted 0xd6c50b442de47f70
(config-keychain)#key 1
(config-keychain-key)#no key-string
```

lsp-gen-interval

Use this command to set the minimum interval before regenerating the same link-state packet (LSP). The smaller the interval, the faster the convergence. However, this setting might cause more frequent flooding.

Use the `no` parameter with this command to set the interval to the default.

Command Syntax

```
lsp-gen-interval <1-120>
lsp-gen-interval (level-1|level-2) <1-120>
no lsp-gen-interval (level-1|level-2|)
```

Parameters

`<1-120>`	Minimum interval in seconds.
`level-1`	Interval for Level-1 IS.
`level-2`	Interval for Level-2 IS.

Default

By default, ISIS uses 30 seconds for the interval and the interval is applied to both level-1 and level-2.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router isis bb
(config-router)#lsp-gen-interval 5
```

lsp-mtu

Use this command to set maximum transfer unit for link-state packets (LSPs).

Use the `no` parameter with this command to set the interval to the default.

Command Syntax

```
lsp-mtu (level-1|level-2|) <512-4352>
no lsp-mtu (level-1|level-2|)
```

Parameters

`<512-4352>`	MTU size
`level-1`	Size for Level-1 IS.
`level-2`	Size for Level-2 IS.

Default

By default, the MTU is 1492 bytes.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router isis bb
(config-router)#lsp-mtu 555
```

lsp-refresh-interval

Use this command to set the link-state packet (LSP) refresh interval.

IP Infusion Inc. recommends making the `lsp-refresh-interval` smaller than the max-lsp-lifetime value.

Use the `no` parameter to set the interval to the default value.

Command Syntax

```
lsp-refresh-interval <1-65535>
no lsp-refresh-interval
```

Parameters

`<1-65535>`	Refresh interval in seconds.

Default

By default, the interval is 900 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#lsp-refresh-interval 600

(config)#router isis bb
(config-router)#no lsp-refresh-interval
```

max-area-address

Use this command to set the maximum number of ISIS areas that can be configured on this router with the `net` command. By default, ISIS permits a maximum of three areas that can be defined on a router.

Use the `no` parameter with this command to set the maximum number of ISIS areas to its default (3).

Command Syntax

```
max-area-address <3-254>
no max-area-address
```

Parameters

`<3-254>`	The maximum number of areas in the network.

Default

By default, the maximum number of areas is 3.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis net2
(config-router)#max-area-address 5
(config-router)#net 71.0001.0000.0000.0001.00
(config-router)#net 72.0001.0000.0000.0001.00
(config-router)#net 73.0001.0000.0000.0001.00
(config-router)#net 74.0001.0000.0000.0001.00
(config-router)#net 75.0001.0000.0000.0001.00
```

© 2025 IP Infusion Inc. Proprietary

max-lsp-lifetime

Use this command to set the maximum link-state packet (LSP) lifetime. You must set `max-lsp-lifetime` greater than the lsp-refresh-interval interval.

Use the `no` parameter to set the lifetime to the default.

Command Syntax

```
max-lsp-lifetime <350-65535>
no max-lsp-lifetime
```

Parameters

`<350-65535>`	Maximum LSP lifetime in seconds.

Default

By default, max-lsp-lifetime is set to 1200 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#max-lsp-lifetime 1500

(config)#router isis bb
(config-router)#no max-lsp-lifetime
```

© 2025 IP Infusion Inc. Proprietary

metric-style

Use this command to configure the ISIS metric style. Use the following table when changing the method of how TLV encodes and SPF calculates a decision:

Metric-style Command	Wide SPF	Wide TLV	Narrow SPF	Narrow TLV
narrow (default)	OFF	OFF	ON	OFF
narrow transition	ON	OFF	ON	ON
wide	ON	ON	OFF	OFF
wide transition	ON	ON	ON	OFF
transition	ON	ON	ON	ON

Where:

- Wide SPF: Uses wide TLVs for SPF calculation.
- Wide TLV: Encodes wide TLVs in the LSP.
- Narrow SPF: Uses narrow TLVs for SPF calculation.
- Narrow TLV: Encodes narrow TLVs in the LSP.

Use the `no` parameter to set the style to the default style, narrow.

Command Syntax

```
metric-style (narrow|wide|transition) (level-1|level-2|)
metric-style (narrow|wide) transition (level-1|level-2|)
no metric-style
```

Parameters

`narrow`	Specify the old style of TLVs with narrow metric.
`wide`	Specify the new style of TLVs to carry wider metric.
`transition`	Specify to send and accept both styles of TLVs during transition.
`level-1`	Specify the level-1 metric style.
`level-2`	Specify the level-2 metric style.
`transition`	Accept both styles of TLVs during transition

Default

By default, ISIS uses narrow metric style for level 1 and 2.

Command Mode

Router mode

Examples

```
(config)#router isis bb
(config-router)#metric-style narrow transition
```

net

Use this command to add a Network Entity Title (NET) for the instance.

On a router running ISIS, a NET can be 8 to 20 bytes in length. The last byte is always the n-selector, and must be zero. The n-selector 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).

The bytes preceding the system ID are the area ID, which can be from 1 - 13 bytes in length. By default, a maximum of three NETs per router are allowed with a different area ID but the system ID should be the same for all NETs. You can increase the number of area IDs per system ID with the `max-area-address` command.

Use the `no` parameter to remove the NET.

Command Syntax

```
net NET
no net NET
```

Parameters

`NET`	Specify a network entity title (NET) in 1 to 13 octets (that is, XX.XXXX. ... .XXXX.XX).

Default

By default, ISIS does not configure a NET and routing is not enabled for the interface.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router isis bb
(config-router)#net 49.0000.0001.0002.0003.00
```

passive-interface

Use this command to suppress routing updates on all interfaces or on a specified interface, which puts the interfaces into passive mode.

To advertise passive prefixes in LSP, an interface must be configured with ip router isis when using "passive-interface" command. If interface is not configured with isis instance "passive-interface IFNAME" command must be used to add that interface as passive. Enabling passive interface on an ISIS enabled interface disables ISIS routing updates on the interface and makes the interface passive.

Use the `no` parameter with this command to remove interfaces from passive mode

Note: The command `passive-interface IFNAME disable` is used to make a particular interface active after the `passive-interface` command (with no parameters) has set all interfaces as passive.

Command Syntax

```
passive-interface
passive-interface IFNAME
passive-interface IFNAME disable


no passive-interface
no passive-interface IFNAME
no passive-interface IFNAME disable
```

Parameters

`IFNAME`	Interface name. If this parameter is omitted, then all interfaces become passive or active.
`disable`	Disable passive interface.

Command Mode

Router mode

Examples

The following suppresses routing updates on a specified interface.

```
#configure terminal
(config)#router isis 100
(config-router)#passive-interface
(config-router)#passive-interface eth0 disable
```

prc-interval-exp

Use this command to configure exponential back-off delay between PRC calculations.

Use the `no` parameter to disable any set exponential back-off delay between PRC calculations.

Command Syntax

```
prc-interval-exp
prc-interval-exp <0-2147483647> <0-2147483647>
no prc-interval-exp
```

Parameters

`<0-2147483647>`	Set the minimum delay between receiving a change to PRC calculation in milliseconds.
`<0-2147483647>`	Set the maximum delay between receiving a change to PRC calculation in milliseconds.

Default

By default, minimum delay is 500 milliseconds and maximum delay is 50 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis
(config-router)#prc-interval-exp 100 10000

(config)#router isis
(config-router)#no prc-interval-exp
```

redistribute

Use this command to redistribute routes from another protocol into the ISIS routing table.

Use the `no` parameter to disable this function.

Command Syntax

```
redistribute (connected|static|rip|ospf|bgp) {metric <0-16777215>| metric-type
  (internal|external)|level-1|level-2|level-1-2|route-map WORD}
no redistribute (connected|static|rip|ospf|bgp)  {metric <0-16777215>| metric-type
  (internal|external)|level-1|level-2|level-1-2|route-map WORD}
```

Parameters

`connected`	Redistribute connected routes.
`static`	Redistribute static routes.
`rip`	Redistribute RIP routes.
`ospf`	Redistribute OSPF routes.
`bgp`	Redistribute BGP routes.
`metric`	Specify the metric for redistributed routes.
`<0-16777215>`	Specify the IS-IS default metric.
`metric-type`	Specify the IS-IS exterior metric type for redistributed routes:
`internal`	Set IS-IS internal metric type.
`external`	Set IS-IS external metric type.
`level-1`	Redistribute routes into level 1 only
`level-2`	Redistribute routes into level 2 only (default)
`level-1-2`	Redistribute routes into both levels.
`route-map`	Specify a Route map reference.
`WORD`	Specify name of the route-map.

Default

By default, redistribute command is disabled.

If no level parameter is specified, by default redistributed routes will be added in level-2 LSP. If is-type of an IS-IS instance is level-1, the level parameter must be set to level-1 for redistribute command to take effect and redistribute routes in L1 LSP.

Command Mode

For Ipv4: Router mode

For Ipv6: address-family ipv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
>ena
#con term
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#router isis A
(config-router)#redistribute bgp metric 12
```

redistribute isis

Use this command to redistribute reachability information from one level to the other level. If an distribute-list name is given with this command for an access list that does not exist, the routes are still redistributed.

Use the `no` parameter with this command to stop redistribution.

Command Syntax

```
redistribute isis level-1 into level-2
redistribute isis level-2 into level-1
redistribute isis level-1 into level-2 distribute-list WORD
redistribute isis level-2 into level-1 distribute-list WORD
no redistribute isis level-1 into level-2
no redistribute isis level-2 into level-1
```

Parameters

`level-1`	Specify an inter-area route from level-1.
`level-2`	Specify an inter-area routes from level-2.
`into`	Specify a level from level-n into level-m.
`level-1`	Specify an inter-area route into level-1.
`level-2`	Specify an inter-area routes into level-2.
`distribute-list`	Indicate the distributed-list parameter.
`WORD`	Specify the actual selected route.

Default

By default, ISIS redistributes selected L1 routes into L2.

Command Mode

For Ipv4: Router mode

For Ipv6: address-family ipv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#redistribute isis level-2 into level-1

(config)#router isis bb
(config-router)#redistribute isis level-2 into level-1 distribute-list new

(config)#router isis bb
(config-router)#no redistribute isis level-2 into level-1
```

redistribute isis WORD

Use this command to redistribute reachability information from one isis instance into another instance. Only one isis instance can be redistributed.

Use the `no` parameter with this command to stop redistribution.

Note: This command is not supported for address family IPv6.

Command Syntax

```
redistribute isis WORD {metric <0-16777215>| metric-type (internal|external)
  |level-1|level-2|level-1-2|route-map WORD}
```

```
no redistribute isis WORD {metric <0-16777215>| metric-type
  (internal|external)|level-1|level-2|level-1-2|route-map WORD}
```

Parameters

`WORD`	Specify an isis instance name or tag
`metric`	Specify the metric for redistributed routes.
`<0-16777215>`	Specify the IS-IS default metric.
`metric-type`	Specify the IS-IS exterior metric type for redistributed routes:
`internal`	Set IS-IS internal metric type.
`external`	Set IS-IS external metric type.
`level-1`	IS-IS Level-1 routes.
`level-2`	IS-IS Level-2 routes.
`level-1-2`	IS-IS Level-1 and Level-2 routes.
`route-map`	Specify a Route map reference.
`WORD`	Specify name of the route-map.

Default

By default, redistribute command is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
>ena
#con term
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#router isis A
(config-router)#redistribute isis B metric 12
```

router isis

Use this command to initiate an ISIS routing instance. This command starts the ISIS routing instance and enters router configuration mode. Configure at least one NET to start routing. Also, enable a particular interface with the ip router isis command.

Use the `no` parameter with this command to remove an ISIS routing instance.

Command Syntax

```
router isis WORD
no router isis WORD
```

Parameters

`WORD`	Name that specifies an ISO routing instance tag.

Default

By default, ISIS routing instance is not configured.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router isis New
(config-router)#
```

send-lifetime

Use this command to specify the time period during which the authentication key on a key chain can be sent.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> HH:MM:SS MONTH <01-31> <1993-2035>
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> infinite
send-lifetime HH:MM:SS MONTH <01-31> <1993-2035> duration <1-2147483646>
no send-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of send-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to start. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to start as the first three letters of the month with first letter in caps, for example, Jan. (case sensitive)
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the end time of send-lifetime in hours, minutes and seconds.
`<01-31>`	Specify the day of the month to end. If the day is a single-digit, the leading 0 must be added, example: 01, 02, 03, etc.
`MONTH`	Specify the month of the year to end as the first three letters of the month with first letter in caps, for example, Jan. (case sensitive)
`<1993-2035>`	Specify the year to end.
`duration`	Indicate the duration parameter.
`<1-2147483646>`	Specify the actual end time duration of a key in seconds.
`infinite`	Specify the end time to never expire.

Applicability

No default value is specified

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of `send-lifetime` for `key 1` on the key chain named `mychain`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
```

```
(config-keychain-key)#send-lifetime 03:03:01 Jan 03 2004 04:04:02 Dec 06 2006
```

set-overload-bit

Use this command to set the overload-bit in self link-state packets (LSPs), preventing the router from being used as a transit router during SPF calculation. The router will still receive LSPs even when the overload bit is set.

Use the no parameter to clear the overload-bit from self-LSPs.

Command Syntax

```
set-overload-bit ({suppress (external|interlevel|external interlevel|interlevel
  external)|on-startup (<5-86400>|wait-for-bgp)}|)
no set-overload-bit
```

Parameters

`suppress`	Suppresses the redistribution of the specified types of reachability data during overload state.
`external`	Suppresses the redistribution of external reachability, preventing the advertisement of IP prefixes learned from other protocols.
`interlevel`	Suppresses the redistribution of interlevel reachability.
`external interlevel` or `interlevel external`	Suppresses the redistribution of both external and interlevel reachability.
`on-startup`	Sets the overload bit only at startup and clears the bit after the specified interval has elapsed.
`<5-86400>`	Specifies the time in seconds to advertise oneself as overloaded after a reboot.
`wait-for-bgp`	Sets the overload bit at startup, then clears the bit after BGP signals it has finished converging, or clears it after 10 minutes if BGP does not signal convergence. If BGP is not running, clears the overload bit immediately.

Default

The overload-bit is not set by default.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.4.2.

Example

This example configures the router to set the overload bit at startup, does not allow the unset of the overload bit until BGP converges, suppresses redistribution between levels, and suppresses redistribution from external routing protocols while the overload bit is set.

```
OcNOS(config)#router isis bb
OcNOS(config-router)#set-overload-bit on-startup wait-for-bgp suppress
interlevel external
```

snmp restart isis

Use this command to restart SNMP in Intermediate System to Intermediate System (IS-IS)

Command Syntax

```
snmp restart isis
```

Parameters

None

Default

By default, snmp restart is not configured.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart isis
```

spf-interval-exp

Use this command to set the minimum and maximum hold intervals between Shortest Path First (SPF) calculations.

The `spf-interval-exp` command configures the minimum and maximum interval time between the receipt of a topology change and the calculation of the SPF.

Use the `no` parameter with this command to set the minimum and maximum hold intervals to the default.

Command Syntax

```
spf-interval-exp <0-2147483647> <0-2147483647>
spf-interval-exp (level-1|level-2) <0-2147483647> <0-2147483647>
no spf-interval-exp ((level-1|level-2)|)
```

Parameters

`<0-2147483647>`	Specify the minimum delay between receiving a change to the SPF calculation in milliseconds. The default SPF minimum hold-time interval value is 500 milliseconds.
`<0-2147483647>`	Specify the maximum delay between receiving a change to the SPF calculation in milliseconds. The default SPF maximum hold-time interval value is 50 seconds.
`level-1`	Specify an interval for Level-1 IS.
`level-2`	Specify an interval for Level-2 IS.

Default

By default, ISIS uses 500 milliseconds and 50,000 milliseconds for the minimum and maximum hold intervals, respectively. The values are applied to both level-1 and level-2 if the `level` parameter is omitted.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#spf-interval-exp level-1 600 60000

(config)#router isis bb
(config-router)#no spf-interval-exp level-1
```

summary-address

Use this command to configure Summary Address to summarize IPv4 reachability information.

Use the `no` parameter with this command to unconfigure the summary.

Command Syntax

```
summary-address A.B.C.D/M (level-1 |level-1-2 |level-2) (metric <1-63>|)
no summary-address A.B.C.D/M
```

Parameters

`A.B.C.D/M`	Specify the IPv4 prefix to be announced.
`level-1`	Specify the reachability information only for Level-1.
`level-1-2`	Specify the reachability information for both Level-1 and Level-2.
`level-2`	Specify the reachability information only for Level-2.
`metric`	Specify the metric for the summarized address.
`<1-63>`	Specify the metric. The default is 0.

Default

By default, ISIS does not configure the summary-address. Level must be configured along with summary-address. Metric is optional.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis bb
(config-router)#summary-address 10.10.0.0/16 level-1-2 metric 3

(config)#router isis bb
(config-router)#no summary-address 10.10.0.0/16
```

summary-prefix

Use this command to configure Summary Prefix to summarize IPv6 reachability information.

Use the `no` parameter with this command to unconfigure the summary.

Command Syntax

```
summary-prefix X:X::X:X/M (level-1 |level-1-2 |level-2) (metric <1-63>|)
no summary-prefix X:X::X:X/M
```

Parameters

`X:X::X:X/M`	Specify the IPv6 prefix to be announced.
`level-1`	Specify the reachability information only for Level-1.
`level-1-2`	Specify the reachability information for both Level-1 and Level-2.
`level-2`	Specify the reachability information only for Level-2.
`metric`	Specify the metric for the summarized address.
`<1-63>`	Specify the metric. The default is 0.

Default

By default, ISIS does not configure the summary-prefix. Level must be configured along with summary-prefix. Metric is optional.

Command Mode

Address Family IPv6 mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#address-family ipv6
(config-router-af)#summary-prefix 2356::1/64 level-1 metric 30
```

CHAPTER 2 IS-IS Graceful Restart Commands

This chapter describes the IS-IS graceful restart commands:

- capability restart graceful
- isis restart grace-period
- isis restart-hello-interval
- isis restart helper
- isis restart suppress-adjacency
- restart isis graceful
- restart-timer

capability restart graceful

Use this command to enable the graceful restart capability.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
capability restart graceful
no capability restart graceful
```

Parameters

NA

Default

By default, graceful restart capability is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example enables the graceful restart capability on a router.

```
#configure terminal
(config)#router isis bb
(config-router)#capability restart graceful
```

isis restart grace-period

Use this command to configure the T3 timer, the time the restarting router retains the forwarding table.

Use the `no` parameter to use the default value.

Command Syntax

```
isis restart grace-period <1-65535>
no isis restart grace-period
```

Parameters

`<1-65535>`	Grace period in seconds.

Default

By default, ISIS uses 65535 for the period value, and the value is applied to both level-1 and level-2.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example enables and then disables a restart grace period at one second.

```
#configure terminal
(config)#isis restart grace-period 1

(config)#no isis restart grace-period
```

isis restart-hello-interval

Use this command to configure the T1 timer, interval of ISIS Hello packet with restart TLV.

Use the `no` parameter to use the default value.

Command Syntax

```
isis restart-hello-interval <1-65535> (level-1|level-2|)
no isis restart-hello-interval (level-1|level-2|)
```

Parameters

`<1-65535>`	Specify the number of seconds in the interval.
`level-1`	Specify the hello-interval for level-1 IIHs.
`level-2`	Specify the hello-interval for level-1 IIHs.

Default

By default, ISIS uses 3 seconds for the hello value, and the interval is applied to both level-1 and level-2.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example enables and then disables a restart hello interval at 123 seconds for a level 1 interface.

```
#configure terminal
(config)#interface eth0
(config-if)#isis restart-hello-interval 123 level-1

(config-if)#no isis restart-hello-interval level-1
```

isis restart helper

Use this command to configure the router's helper mode capability.

Use the `no` parameter to disable the helper mode for this router.

Command Syntax

```
isis restart helper
no isis restart helper
```

Parameters

None

Default

By default, most routers are not a restart helper router.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example enables and then disables ISIS restart helper.

```
#configure terminal
(config)#isis restart helper

(config)#no isis restart helper
```

isis restart suppress-adjacency

Use this command to enable ISIS to request that its adjacency be suppressed after the ISIS daemon process starts or restarts until the Link State Packet Database (LSPDB) synchronizes.

Use the `no` parameter to disable suppress-adjacency.

Command Syntax

```
isis restart suppress-adjacency
no isis restart suppress-adjacency
```

Parameters

None

Default

By default, ISIS does not request that its adjacency be suppressed after the ISIS daemon process starts or restarts.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example enables and then disables ISIS restart suppress adjacency.

```
#configure terminal
(config)#isis restart suppress-adjacency

(config)#no isis restart suppress-adjacency
```

restart isis graceful

Use this command to restart the ISIS router.

Command Syntax

```
restart isis graceful (grace-period <1-65535>|)
```

Parameters

<1-65535>	Grace period in seconds.

Default

By default, the ISIS router is not restarted gracefully. Default grace-period is 65535 seconds.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#restart isis graceful grace-period 60

#restart isis graceful
% Warning : ISISD process will stop and needs to restart manually,
You may loose ISIS configuration, if not saved
Proceed for graceful restart? (y/n):y
```

restart-timer

Use this command to restart the ISIS T2 timer. When a node comes up after graceful restart, within this time, the LSPDB sync should be completed.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
restart-timer <5-65535> (level-1|level-2|)
no restart-timer (level-1|level-2|)
```

Parameters

`<5-65535>`	Restart time in seconds
`level-1`	Restart is only for Level-1.
`level-2`	Restart is only for Level-2.

Default

The default value is 60 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example enables and then disables the restart timer at 555 seconds for a level 2 interface.

```
#configure terminal
(config)#router isis bb
(config-router)#restart-timer 555 level-2

(config)#router isis bb
(config-router)#no restart-timer level-2
```

CHAPTER 3 IS-IS Show Commands

This chapter provides describes the IS-IS show commands:

- show clns is-neighbors
- show clns neighbors
- show debugging isis
- show ip isis igp-shortcut-lsp
- show ip isis route
- show ip isis route igp-shortcut
- show ip protocols
- show ip route fast-reroute
- show ip isis route fast-reroute
- show ip isis lfa-config
- show isis counter
- show isis database
- show isis interface
- show isis spf-logs
- show isis topology
- show running-config interface isis
- show running-config router isis

show clns is-neighbors

Use this command to display IS neighbor adjacencies.

Command Syntax

```
show clns is-neighbors (detail|)
show clns WORD is-neighbors (detail|)
show clns is-neighbors IFNAME (detail|)
show clns WORD is-neighbors IFNAME (detail|)
```

Parameters

`detail`	Detailed information.
`WORD`	Information for a single IS-IS area.
`IFNAME`	Information for a single interface.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show clns is-neighbors detail
Tag abc: VRF : default
System Id      Interface   State  Type Priority  Circuit Id
0000.0000.0003 eth1        Up     L1   64        0000.0000.0003.01
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 00:12:31
Area Address(es): 52
IP Address(es): 11.11.11.2
Level-1 Protocols Supported: IPv4
Adjacency advertisement: Advertise
#
```

Table 3-1 explains the fields in the output.

Table 3-1: show clns is-neighbors output

Field	Description
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
System Id	Uniquely identifies a system within an area.
Interface	Interface from which the system was learned.

Table 3-1: show clns is-neighbors output (Continued)

Field	Description
State	Adjacency state: Init: Router is an IS and is waiting for an IS-IS hello message. IS-IS regards the neighbor as not adjacent. Up: IS is considered reachable
Type	Type of adjacency: L1: Router adjacency for level 1 routing only L2: Router adjacency for level 2 only L1L2: Router adjacency for level 1 and level 2 routing
Priority	IS-IS priority that the respective neighbor is advertising. The highest priority neighbor is elected the designated IS-IS router for the interface.
Circuit Id	Number that the router uses to uniquely identify its IS-IS interface. When the interface is attached to a broadcast network, the Circuit ID is concatenated with System ID of the designated router for the interface.
Adjacency ID	Adjacency identifier.
Uptime	How long the adjacency has existed.
Area Addresses	Area addresses associated with the intermediate-system adjacencies.
IPv4/IPv6 address(es)	IP addresses of the ES or IS.
Protocols Supported	IPv4 and/or IPv6.
Adjacency advertisement	Restart: Suppress or Advertise.

show clns neighbors

Use this command to display ES and IS neighbor adjacencies.

Command Syntax

```
show clns neighbors (detail|)
show clns WORD neighbors (detail|)
show clns neighbors IFNAME (detail|)
show clns WORD neighbors IFNAME (detail|)
```

Parameters

`detail`	Detailed information for all interfaces.
`WORD`	Information for a single IS-IS area.
`IFNAME`	Information for a single interface.

Command Mode

Exec mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show clns neighbors detail

Tag abc: VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0003 eth1        0800.277b.411d  Up     6         L1   IS-IS
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 00:15:58
Area Address(es): 52
IP Address(es): 11.11.11.2
Level-1 Protocols Supported: IPv4
Adjacency advertisement: Advertise
#
```

Table 3-2 explains the fields in the output.

Table 3-2: show clns neighbors output

Field	Description
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
System Id	Uniquely identifies a system within an area.
Interface	Interface from which the system was learned.

Table 3-2: show clns neighbors output (Continued)

Field	Description
SNPA	SubNetwork Point of Attachment (SNPA): MAC address of the next-hop.
State	Adjacency state: Init: Router is an IS and is waiting for an IS-IS hello message. IS-IS regards the neighbor as not adjacent. Up: ES or IS is considered reachable
Holdtime	Number of seconds before this adjacency entry times out.
Type	Type of adjacency: L1: Router adjacency for level 1 routing only L2: Router adjacency for level 2 only L1L2: Router adjacency for level 1 and level 2 routing
Protocol	Protocol through which the adjacency was learned.
Adjacency ID	Adjacency identifier.
Uptime	How long the adjacency has existed.
Area Addresses	Area addresses associated with the intermediate-system adjacencies.
IPv4/IPv6 address(es)	IP addresses of the ES or IS.
Topology	IPv4 and/or IPv6.
Protocols Supported	IPv4 and/or IPv6.
Adjacency advertisement	Restart: Suppress or Advertise.

show debugging isis

Use this command to display the status of the debugging of the ISIS system.

Command Syntax

```
show debugging isis
```

Parameters

None

Command Mode

Exec mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging isis
IS-IS debugging status:
IS-IS Interface FSM debugging is on
IS-IS Neighbor FSM debugging is on
IS-IS events debugging is on
IS-IS PDU debugging is on
IS-IS lsp debugging is on
IS-IS spf debugging is on
IS-IS NSM debugging is on
IS-IS Check-sum debugging is on
IS-IS Authentication debugging is on
IS-IS Protocol-error debugging is on
IS-IS Local Updates debugging is on
IS-IS Hello debugging is on
IS-IS BFD debugging is on
IS-IS RIB debugging is on
#
```

© 2025 IP Infusion Inc. Proprietary

show ip isis igp-shortcut-lsp

Use this command to display IS-IS shortcut label-switched paths (LSPs).

Command Syntax

```
show ip isis (WORD|) igp-shortcut-lsp
```

Parameters

WORD	Information for a single IS-IS area.

Command Mode

Exec mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip isis igp-shortcut-lsp
#
```

Table 3-3 explains the fields in the output.

Table 3-3: show ip isis igp-shortcut-lsp output

Field	Description
Tunnel-endpoint	Tunnel endpoint address.
Tunnel-id	Tunnel identifier.
Tunnel-metric	Tunnel metric.
active/inactive	Whether the tunnel is active or inactive.

show ip isis route

Use this command to display IS-IS routing table for IPv4.

Command Syntax

```
show ip isis (WORD|) route
```

Parameters

WORD	Information for a single IS-IS area.

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip isis route
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
Tag xyz: VRF : default
          Destination         Metric      Next-Hop            Interface       Tag
L1        10.10.10.0/24       20          11.11.11.1          eth1            0
C         11.11.11.0/24       10          -                   eth1            0
#
```

Header

Each entry in this table has a code preceding it, indicating the source of the routing entry. Table 3-4 shows these codes.

Table 3-4 explains the fields in the output.

Table 3-4: route codes and modifiers

Code	Description
C	Routes directly connected to the local device that were not distributed via IGP. The device inherently knows of these networks, so there is no need to learn about these from another device. Connected routes are preferred over routes for the same network learned from other routing protocols.
E	External.
L1	IS-IS level-1.
L2	IS-IS level-2.
ia	IS-IS inter area (leaked).

Table 3-4: route codes and modifiers (Continued)

Code	Description
D	Discard route. A device performing summarization installs a discard route in its routing table for the summarized network range to prevent routing loops where portions of the summarized network range do not have a more specific route in the RIB. External and internal discard route entries are installed by default. During route summarization, routing loops can happen if data sent to a nonexisting network appears to be a part of the summary, and the router doing the summarization has a less specific route that points back to the sending router for the network.
e	External metric. Routes can be redistributed into IS-IS with either internal or external metrics (internal is the default). The metric type determines the base metric value of the redistributed routes. The value of an internal metric is lower than 64. The value of an external metric is 64-128.

Route Entry Fields

Table 3-5 shows the route entry fields.

Table 3-5: route entry fields

Field	Description
Code	As explained in Table 3-4.
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
Destination	IP address of the remote network.
Metric	ISIS metric used for SPF calculation (1-63). When a route is imported into the IS-IS network without a specified metric, IS-IS uses 10 for the metric value and the value is applied to both level-1 and level-2.
Next-Hop	This route is available through the next hop router located at this IP address. This identifies exactly where packets go when they match this route.
Interface	Interface used to get to the next-hop address for this route.
Tag	Name that identifies the IS-IS area.

show ip isis route igp-shortcut

Use this command to display the IS-IS IGP shortcut routing table.

Command Syntax

```
show ip isis (WORD|) route igp-shortcut
```

Parameters

WORD	Information for an IS-IS area.

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip isis new_isis route igp-shortcut
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, D - discard, e - external metric

Tag aa: VRF : default
     Destination         Metric      Tunnel-ID         Tunnel-End-Point
#
```

Header

Each entry in this table has a code preceding it, indicating the source of the routing entry. Table 3-6 shows these codes.

Table 3-6: Route codes and modifiers

Code	Description
C	Routes directly connected to the local device that were not distributed via IGP. The device inherently knows of these networks, so there is no need to learn about these from another device. Connected routes are preferred over routes for the same network learned from other routing protocols.
E	External.
L1	IS-IS level-1.
L2	IS-IS level-2.
ia	IS-IS inter area (leaked).

Table 3-6: Route codes and modifiers (Continued)

Code	Description
D	Discard route. A device performing summarization installs a discard route in its routing table for the summarized network range to prevent routing loops where portions of the summarized network range do not have a more specific route in the RIB. External and internal discard route entries are installed by default. During route summarization, routing loops can happen if data sent to a nonexisting network appears to be a part of the summary, and the router doing the summarization has a less specific route that points back to the sending router for the network.
e	External metric. Routes can be redistributed into IS-IS with either internal or external metrics (internal is the default). The metric type determines the base metric value of the redistributed routes. The value of an internal metric is lower than 64. The value of an external metric is 64-128.

Route Entry Fields

Table 3-7 shows the route entry fields.

Table 3-7: Route entry fields

Field	Description
Code	As explained in Table 3-4.
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
Destination	IP address.
Metric	Tunnel metric.
Tunnel-ID	Tunnel identifier.
Tunnel-End-Point	Tunnel endpoint address.

show ip protocols

Use this command to display IP process parameters and statistics.

Command Syntax

```
show ip protocols
show ip protocols isis
```

Parameters

None

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip protocols
Routing Protocol is "isis 1 "
Redistributing:
Area Address(es): 52
Distance : (default is 115)
#
```

Table 3-8 explains the output fields.

Table 3-8: show ip protocols output

Field	Description
Routing Protocol	“isis” and the name that identifies the IS-IS area.
Redistributing	Protocols being redistributed (such as RIP, OSPF, or BGP), including metric, metric type (internal or external), and route map.
redistribute isis	Whether redistributing IS-IS level-1 into level-2 and vice versa.
Area Address	Network address of the areas into which the routing process is injecting routes.
Distance: (default is 115)	Administrative distance.

show ip route fast-reroute

Use this command to display Loop Free Alternate Fast Reroute (LFA FRR) routes.

Command Syntax

```
show ip route fast-reroute
```

Parameters

None

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
* - candidate default

IP Route Table for VRF "default"
i L140.40.40.0/24 [115/10] via 10.10.10.142, eth1, 00:00:50
                          [FRR-NH] via 30.30.30.144, eth3

i L150.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:00:50
                          [FRR-NH] via 10.10.10.142, eth1

i L160.60.60.0/24 [115/15] via 10.10.10.142, eth1, 00:00:50
                          [FRR-NH] via 20.20.20.143, eth2
```

show ip isis route fast-reroute

Use this command to display Loop Free Alternate Fast Reroute (LFA FRR) route information and interfaces on which LFA FRR is disabled.

Command Syntax

```
show ip isis (WORD|) route fast-reroute
```

Parameters

`WORD` Routing area tag.

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p - stale info
* - candidate default

IP Route Table for VRF "default"
i L140.40.40.0/24 [115/10] via 10.10.10.142, eth1, 00:00:50
                          [FRR-NH] via 30.30.30.144, eth3

i L150.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:00:50
                          [FRR-NH] via 10.10.10.142, eth1

i L160.60.60.0/24 [115/15] via 10.10.10.142, eth1, 00:00:50
                          [FRR-NH] via 20.20.20.143, eth2
```

show ip isis lfa-config

Use this command to display the Loop Free Alternate Fast Reroute (LFA FRR) tie-break preferences for protection types and the termination hold-on timer.

Command Syntax

```
show ip isis (WORD|) lfa-config (level-1|level-2)
```

Parameters

`WORD` Routing area tag.

`level-1` Level 1 only.

`level-2 Level 2 only`

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#show ip isis lfa-config level-1

TIE-Breaker           Preference values
-------------------------------------------------
Primary Path:                 20
Link Protecting:              30
Node Protecting:              60
Broadcast Interface Disjoint: 70
Secondary Path:                0
Downstream Path:               0
```

Termination Hold On Interval : 1000 ms

show isis counter

Use this command to display the MIB variables used to construct routing tables for IP networks for IS-IS as defined in RFC 4444.

Command Syntax

```
show isis counter
```

Parameters

None

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show isis counter
Tag abc: VRF : default
IS-IS Level-1 isisSystemCounterEntry:
isisSysStatCorrLSPs: 0
isisSysStatAuthTypeFails: 0
isisSysStatAuthFails: 0
isisSysStatLSPDbaseOloads: 0
isisSysStatManAddrDropFromAreas: 0
isisSysStatAttmptToExMaxSeqNums: 0
isisSysStatSeqNumSkips: 0
isisSysStatOwnLSPPurges: 0
isisSysStatIDFieldLenMismatches: 0
isisSysStatMaxAreaAddrMismatches: 0
isisSysStatPartChanges: 0
isisSysStatSPFRuns: 7
isisSysStatPRCRuns: 0

IS-IS Level-2 isisSystemCounterEntry:
isisSysStatCorrLSPs: 0
isisSysStatAuthTypeFails: 0
isisSysStatAuthFails: 0
isisSysStatLSPDbaseOloads: 0
isisSysStatManAddrDropFromAreas: 0
isisSysStatAttmptToExMaxSeqNums: 0
isisSysStatSeqNumSkips: 0
isisSysStatOwnLSPPurges: 0
isisSysStatIDFieldLenMismatches: 0
isisSysStatMaxAreaAddrMismatches: 0
isisSysStatPartChanges: 0
isisSysStatSPFRuns: 3
isisSysStatPRCRuns: 0
#
```

show isis database

Use this command to display link-state database (LSDB) database information. The LSDB is the core of IS-IS routing. All link-state information advertised by neighbors in the same area is stored in the LSDB.

Command Syntax

```
show isis database
show isis database (detail|verbose)
show isis database (detail|verbose) WORD
show isis database (detail|verbose) WORD (l1|l2|level-1|level-2)
show isis database (detail|verbose) (l1|l2|level-1|level-2)
show isis database (detail|verbose) (l1|l2|level-1|level-2) WORD
show isis database WORD
show isis database WORD (l1|l2|level-1|level-2)
show isis database WORD (l1|l2|level-1|level-2) (detail|verbose)
show isis database WORD (detail|verbose)
show isis database WORD (detail|verbose) (l1|l2|level-1|level-2)
show isis database (l1|l2|level-1|level-2)
show isis database (l1|l2|level-1|level-2) (detail|verbose)
show isis database (l1|l2|level-1|level-2) (detail|verbose) WORD
show isis database (l1|l2|level-1|level-2) WORD
show isis database (l1|l2|level-1|level-2) WORD (detail|verbose)
```

Parameters

`detail`	Detailed information.
`verbose`	Verbose information.
`WORD`	Link-state packet (LSP) identifier in the form of XXXX.XXXX.XXXX.XX-XX.
`l1`	IS-IS level-1.
`l2`	IS-IS level-2.
`level-1`	IS-IS level-1.
`level-2`	IS-IS level-2.

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show isis database detail
Area bb:
IS-IS Level-1 Link State Database:
```

```
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
000F.0000.0001.00-00* 0x00000007   0xE15E        1188              1/0/0
  Area Address: 49.000F
  NLPID:        0xCC
  IP Address:   10.10.12.97
  Metric: 10         IP 10.10.12.0 255.255.255.0
  Metric: 10         IS 000F.0000.0001.02
000F.0000.0001.02-00* 0x00000003   0x3C66        1026              1/0/0
  Metric:  0         IS 000F.0000.0001.00
  Metric:  0         IS 000F.0000.0002.00
000F.0000.0002.00-00  0x00000003   0x8C4B        1025              1/0/0
  Area Address: 49.000F
  NLPID:        0xCC
  Hostname:     isisd@redhat
  IP Address:   10.10.12.94
  Metric: 10         IP 10.10.12.0 255.255.255.0
  Metric: 10         IS 000F.0000.0001.02
#
```

Table 3-9 explains the output fields.

Table 3-9: show isis database output

Field	Description
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
IS-IS Level-n Link State Database	Each IS-IS level has a section with Link-State Packet (LSP) information.
LSPID	Link-state packet identifier in the form of XXXX.XXXX.XXXX.XX-XX. The first six octets (“XXXX.XXXX.XXXX”) are the system identifier of the router that originated the LSP. The next octet is the pseudonode identifier: • When this octet is nonzero, the LSP describes links from a designated router (pseudonode) that creates and floods an LSP that describes all systems attached to the network. This mechanism is similar to a router link-state advertisement (LSA) in OSPF. • When this octet is zero, the LSP is from a nonpseudonode which describes the state of the originating router. The last octet is the LSP number. If the value is 0x00, the entire LSP was carried in one LSP. If there is more data than can fit in a single LSP, the LSP is divided into multiple LSP fragments and each fragment has a different LSP number. An asterisk (*) means the LSP originated on the system where the command was given.
LSP Seq Num	LSP sequence number.
LSP Checksum	LSP checksum.
LSP Holdtime	Amount of time the LSP remains valid (in seconds). An LSP hold time of zero means the LSP was purged and is being removed from the link-state database (LSDB) of all routers. The value indicates how long the purged LSP will stay in the LSDB before being completely removed.
ATT	Attached bit. A Level-2 IS indicates its attachment to other areas by setting its attached bit in its Level-1 LSP 0. In other words, this is only set for inter-area routes. Level 1-only routers and Level 1-2 routers that have lost connection to other Level 2 routers will use the attach bit to find the closest Level-2 router. They will point a default route to the closest Level-2 router.

Table 3-9: show isis database output

Field	Description
P	Partition repair. A Level-1 area can become partitioned; this bit means the partition can be repaired via use of Level-2 routes.
OL	Overload bit. Determines whether the IS is congested. When the overload-bit is set in an LSP, other routers will not use this router as a transit router during SPF calculation. Only packets for destinations directly connected to the overloaded router will be sent to this router.

This command also displays information about the IS-IS TLVs in Table 3-10 if present in an LSP. For more about the TLV information, search for “IS-IS TLV Codepoints” on the Interne, check ISO/IEC 10589:2002(E), or other standard mentioned in Table 3-10.

Table 3-10: IS-IS TLV Codepoints

IS-IS TLV Codepoint	Description	Standard
1	Area Addresses	ISO 10589
2	IIS Neighbors	ISO 10589
3	ES Neighbors	ISO 10589
10	Authentication	ISO 10589, RFC 6233
22	Extended IS reachability	RFC 5305
128	IP internal reachability	RFC 1195, RFC 5302
129	Protocols supported	RFC 1195
130	IP external reachability	RFC 1195, RFC 5302
132	IP interface address	RFC 1195
134	Traffic engineering router ID	RFC 5305
135	Extended IP reachability	RFC 5305
137	Host name	RFC 5301, RFC 6233
222	Multi IS reachability	RFC 5120
229	Multi topology	RFC 5120
232	IPv6 interface address	RFC 5308
235	Multi IPv4 reachability	RFC 5120
236	IPv6 reachability	RFC 5308
237	Multi IPv6 reachability	RFC 5120

© 2025 IP Infusion Inc. Proprietary

show isis interface

Use this command to display detailed interface information.

Command Syntax

```
show isis interface
show isis interface IFNAME
show isis interface counter
```

Parameters

IFNAME	Interface name.
counter	Interface counters.

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>show isis interface
eth2 is up, line protocol is up
Routing Protocol: IS-IS (abc)
Network Type: Broadcast
Circuit Type: level-2
Local circuit ID: 0x01
Extended Local circuit ID: 0x00000004
Local SNPA: 0800.2731.a9a0
IP interface address:
10.10.10.1/24
IPv6 interface address:
fe80::a00:27ff:fe31:a9a0/64
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
Number of active level-2 adjacencies: 0
Level-2 LSP MTU: 1492
Next IS-IS LAN Level-2 Hello in 9 seconds
eth1 is up, line protocol is up
Routing Protocol: IS-IS (abc)
Network Type: Broadcast
Circuit Type: level-1
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000003
Local SNPA: 0800.2714.e7f8
IP interface address:
11.11.11.1/24
IPv6 interface address:
fe80::a00:27ff:fe14:e7f8/64
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0003.01
Number of active level-1 adjacencies: 1
Level-1 LSP MTU: 1492
```

```
Next IS-IS LAN Level-1 Hello in 5 seconds

>show isis interface eth1
eth1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000003
Local SNPA: 0800.27e3.0e64
IP interface address:
30.0.0.1/24
IPv6 interface address:
fe80::a00:27ff:fee3:e64/64
Holddown timer = 100 seconds, Remaining time = 90 seconds
Level-1 Metric: 63/16777214, Priority: 64, Circuit ID: 0000.0000.0001.02
Number of active level-1 adjacencies: 1
Level-1 LSP MTU: 1492
Next IS-IS LAN Level-1 Hello in 1 seconds
```

Table 3-11 explains the output fields.

Table 3-11: show isis interface

Field	Description
IFNAME is up, line protocol is up/down	Whether the interface is up or down
Routing Protocol	“IS-IS” and the name that identifies the IS-IS instance.
Network Type	• Broadcast • Point-to-Point • Loopback
Circuit Type	Whether the interface is configured for: • Local routing: level-1 • Area routing: level 2 • Local and area routing: level-1-2
Local circuit ID	Local circuit identifier assigned when interface is created. Each IS-IS interface is assigned a circuit identifier to identify the interface within the link-state database. If the interface is attached to a multiaccess network, the circuit ID is concatenated with the system ID of the designated IS. This is called the pseudonode ID.
Extended Local circuit ID	Interface index.
Local SNPA	SubNetwork Point of Attachment: for broadcast networks, MAC address.
IP interface address	IPv4 addresses assigned to IS-IS interface.
IPv6 interface address	IPv6 addresses assigned to IS-IS interface.
Level-1 Metric	Interface metric value; used for SPF calculation.
Priority	Priority for designated IS election.

Table 3-11: show isis interface (Continued)

Field	Description
Circuit ID	Unique ID assigned to a circuit internally.
Number of active level-1 adjacencies	Number of adjacencies formed with a neighboring router.
Level-1 LSP MTU	Maximum transmission unit: maximum transmission size for a packet on this interface.
Level-2 Metric	Interface metric value; used for SPF calculation.
Priority	Priority for designated IS election.
Circuit ID	Unique ID assigned to a circuit internally.
Number of active level-2 adjacencies	Number of adjacencies formed with a neighboring router.
Level-2 LSP MTU	Maximum transmission unit: maximum transmission size for a packet on this interface.
Next IS-IS LAN Level-1 Hello	For broadcast networks, when the next IS hello will be sent on this interface.
Next IS-IS LAN Level-2 Hello	For broadcast networks, when the next IS hello will be sent on this interface.
Next IS-IS Hello in	For point-to-point networks, when the next IS hello will be sent on this interface.
Bandwidth	Traffic engineering: interface bandwidth.
Maximum reservable bandwidth	Traffic engineering: maximum reservable interface bandwidth.
Available bandwidth at priority	Traffic engineering: available interface bandwidth at priority.
Bidirectional Forwarding Detection is disabled/enabled/ configured	BFD state

show isis spf-logs

Use this command to display the Shortest Path First (SPF) related information. It also display Partial Route Calculation (PRC) related information only when prc-interval-exp is enabled.

Command Syntax

```
show isis (WORD|) spf-logs (level-1|level-2|level-1-2)
```

Parameters

WORD	Routing area tag.
level-1	Level 1 only.
level-2	Level 2 only.
level-1-2	Level 1-2

Command Mode

Exec mode, Privileged exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
OcNOS#show isis spf-logs level-1
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 5 times
 SPF algorithm last executed 00:05:06.106 ago

PRC logs:
Next PRC is not scheduled yet
PRC schedule delay min 16 secs 0 msecs
PRC schedule delay max 65 secs 0 msecs
PRC algorithm executed 2 times
PRC algorithm last executed 00:03:58.256 ago
```

© 2025 IP Infusion Inc. Proprietary

show isis topology

Use this command to display paths to Intermediate Systems.

Command Syntax

```
show isis topology (l1|l2|level-1|level-2|)
show isis WORD topology (l1|l2|level-1|level-2|)
```

Parameters

WORD	Display information for specified instance.
l1	Display the path to all level-1 routers in the area.
l2	Display the path to all level-2 routers in the domain.
level-1	Display the path to all level-1 routers in the area.
level-2	Display the path to all level-2 routers in the domain.

Command Mode

Exec mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show isis topology

Tag abc: VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop            Interface   SNPA
000F.0000.0001       --
000F.0000.0002       10         000F.0000.0002      eth2        0006.5B0E.D27D

IS-IS paths to level-2 routers
System Id            Metric     Next-Hop            Interface   SNPA
0000.0000.0001       10         0000.0000.0001      eth2        0000.0CFA.F002
```

Table 3-12 explains the output fields.

Table 3-12: show isis topology output

Field	Description
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
IS-IS paths to level-n routers	Each IS-IS level has a section with topology information.
System Id	Uniquely identifies a system within an area.
Metric	ISIS metric used for SPF calculation (1-63).

Table 3-12: show isis topology output

Field	Description
Next-Hop	This route is available through the next hop router located at this IP address.
Interface	Interface from which the system was learned.
SNPA	SubNetwork Point of Attachment (SNPA): MAC address of the device.

show running-config interface isis

Use this command to display the ISIS interface configuration.

Command Syntax

```
show running-config interface IFNAME isis
```

Parameters

IFNAME	Interface name.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show running-config interface eth0 isis
!
interface eth0
 isis tag 500 level-1
!
```

show running-config router isis

Use this command to display the ISIS router configuration.

Command Syntax

```
show running-config router isis
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config-router)#show running-config router isis
!
router isis 1
passive-interface eth1
!
```

Open Shortest Path First Configuration

CHAPTER 1 OSPFv2

This chapter contains basic OSPFv2 (Open Shortest Path First) configuration examples.

Enable OSPF on an Interface

The diagram shows the minimum configuration required to enable OSPF on an interface. R1 and R2 are two routers in Area 0 connecting to network 10.10.10.0/24.

Note: Configure one interface so that it belongs to only one area. It is possible, however, to configure different interfaces on a router to belong to different areas.

Topology

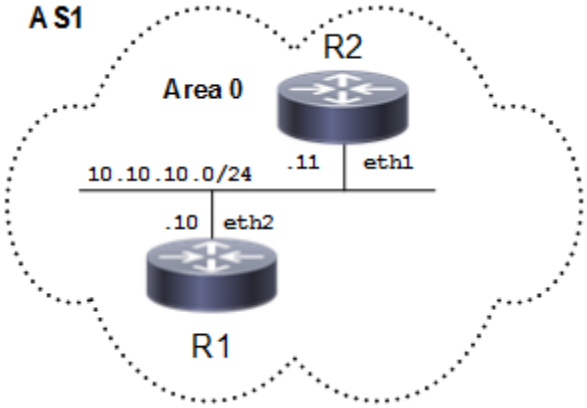

Figure 1-28: Basic OSPF Topology

R1

`#configure terminal`	Enter configure mode
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).

R2

`#configure terminal`	Enter configure mode
`(config)#router ospf 200`	Configure the routing process, and specify the Process ID `(200)`. The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface `(10.10.10.0/24)` on which OSPF runs, and associate the area ID `(0)` with the interface.

Validation

R1

```
#show ip ospf
 Routing Process "ospf 100" with ID 10.12.26.88
 Process uptime is 1 minute
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incomming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 4
 Number of LSA received 4
 Number of areas attached to this router: 1
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:00:08.102 ago
        SPF algorithm executed 3 times
        Number of LSA 7. Checksum 0x0312b5
 Dste Staus: Disabled

#show ip ospf interface
eth2 is up, line protocol is up
  Internet Address 10.10.10.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.88, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Backup Designated Router (ID) 10.12.26.89, Interface Address 10.10.10.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:11
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 13 sent 19, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 3 sent 5
  LS-Ack received 3 sent 3, Discarded 0
No authentication
```

© 2025 IP Infusion Inc. Proprietary

```
#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.89       1   Full/Backup     00:00:39    10.10.10.11     eth2
0


#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

C  10.10.10.0/24 [1] is directly connected, eth2, Area 0.0.0.0
```

R2

```
#show  ip ospf
 Routing Process "ospf 200" with ID 10.12.26.89
 Process uptime is 1 minute
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incomming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 3
 Number of LSA received 5
 Number of areas attached to this router: 1
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:00:45.638 ago
        SPF algorithm executed 4 times
        Number of LSA 7. Checksum 0x0312b5
 Dste Staus: Disabled

#show ip ospf interface
eth1 is up, line protocol is up
  Internet Address 10.10.10.11/24, Area 0.0.0.0, MTU 1500
```

```
  Process ID 200, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Backup Designated Router (ID) 10.12.26.89, Interface Address 10.10.10.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 30 sent 31, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 5 sent 3
  LS-Ack received 2 sent 3, Discarded 0
No authentication


#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 200 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/DR         00:00:33    10.10.10.10     eth1
0

#show ip ospf route

OSPF process 200:
Codes: C - connected, D - Discard, 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

C  10.10.10.0/24 [1] is directly connected, eth1, Area 0.0.0.0
```

Set Priority

This example shows how to set the priority for an interface. Set a high priority for a router to make it the Designated Router (DR). Router R3 is configured to have a priority of 10, which is higher than the default priority (1) of R1 and R2; making it the DR.

Topology

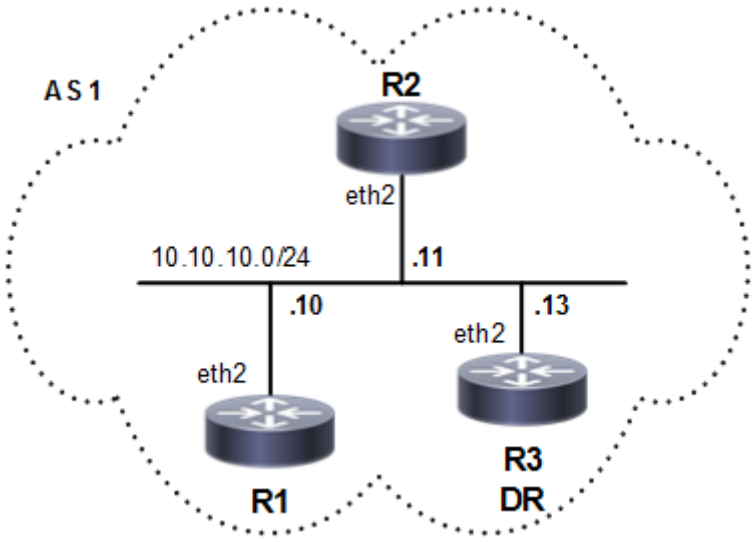

Figure 1-29: Set OSPF Priority

R3

`#configure terminal`	Enter configure mode
`(config)#``interface eth2`	Enter interface mode.
`(config-if)#``ip ospf priority 10`	Specify the router priority to a higher priority (10) to make R3 the Designated Router (DR).
`(config-if)#``exit`	Exit interface mode.
`(config)#``router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#``network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface.

R1

`#configure terminal`	Enter configure mode
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).

R2

`#configure terminal`	Enter configure mode
`(config)#router ospf 200`	Configure the routing process, and specify the Process ID (`200`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface.

Validation

R1

```
#sh ip ospf neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.89       1   Full/DROther    00:00:39    10.10.10.11     eth2
0
10.12.26.90      10   Full/DR         00:00:32    10.10.10.13     eth2
0

#sh ip ospf interface
eth2 is up, line protocol is up
  Internet Address 10.10.10.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.88, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.90, Interface Address 10.10.10.13
  Backup Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:10
  Neighbor Count is 2, Adjacent neighbor count is 2
  Suppress hello for 0 neighbor(s)
  Hello received 30 sent 19, DD received 6 sent 8
  LS-Req received 2 sent 2, LS-Upd received 16 sent 6
  LS-Ack received 8 sent 7, Discarded 0
  No authentication

#sh running-config
!
no service password-encryption
!
hostname rtr1
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
```

```
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.88/24
!
interface eth1
!
interface eth2
 ip address 10.10.10.10/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
interface eth7
!
router ospf 100
 network 10.10.10.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end
```

R2

```
#sh running-config
!
no service password-encryption
!
hostname R2
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.89/24
```

```
!
interface eth1
!
interface eth2
 ip address 10.10.10.11/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
router ospf 200
 network 10.10.10.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

#sh ip ospf neighbor

Total number of full neighbors: 2
OSPF process 200 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/Backup     00:00:30    10.10.10.10     eth2
0
10.12.26.90      10   Full/DR         00:00:31    10.10.10.13     eth2
0
R2#

R2#sh ip ospf interface
eth2 is up, line protocol is up
  Internet Address 10.10.10.11/24, Area 0.0.0.0, MTU 1500
  Process ID 200, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DROther, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.90, Interface Address 10.10.10.13
  Backup Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
  Neighbor Count is 2, Adjacent neighbor count is 2
  Suppress hello for 0 neighbor(s)
  Hello received 71 sent 36, DD received 7 sent 7
  LS-Req received 2 sent 2, LS-Upd received 9 sent 4
  LS-Ack received 3 sent 4, Discarded 1
  No authentication
```

R3

```
#sh running-config
!
```

```
no service password-encryption
!
hostname R3
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
spanning-tree mode provider-rstp

ethernet cfm enable
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.90/24
!
interface eth1
!
interface eth2
 ip address 10.10.10.13/24
 ip ospf priority 10
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
router ospf 100
 network 10.10.10.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

#sh ip ospf neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/Backup     00:00:33    10.10.10.10     eth2
0
10.12.26.89       1   Full/DROther    00:00:30    10.10.10.11     eth2
0
```

```
#sh ip ospf interface
eth2 is up, line protocol is up
  Internet Address 10.10.10.13/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.90, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 10, TE Metric 1
  Designated Router (ID) 10.12.26.90, Interface Address 10.10.10.13
  Backup Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Neighbor Count is 2, Adjacent neighbor count is 2
  Suppress hello for 0 neighbor(s)
  Hello received 99 sent 60, DD received 8 sent 6
  LS-Req received 2 sent 2, LS-Upd received 9 sent 12
  LS-Ack received 9 sent 6, Discarded 1
```

No authentication

Area Border Router

This example shows configuration for an Area Border Router. R2 is an Area Border Router (ABR). On R2, Interface eth0 is in Area 0, and Interface eth1 is in Area 1.

Topology

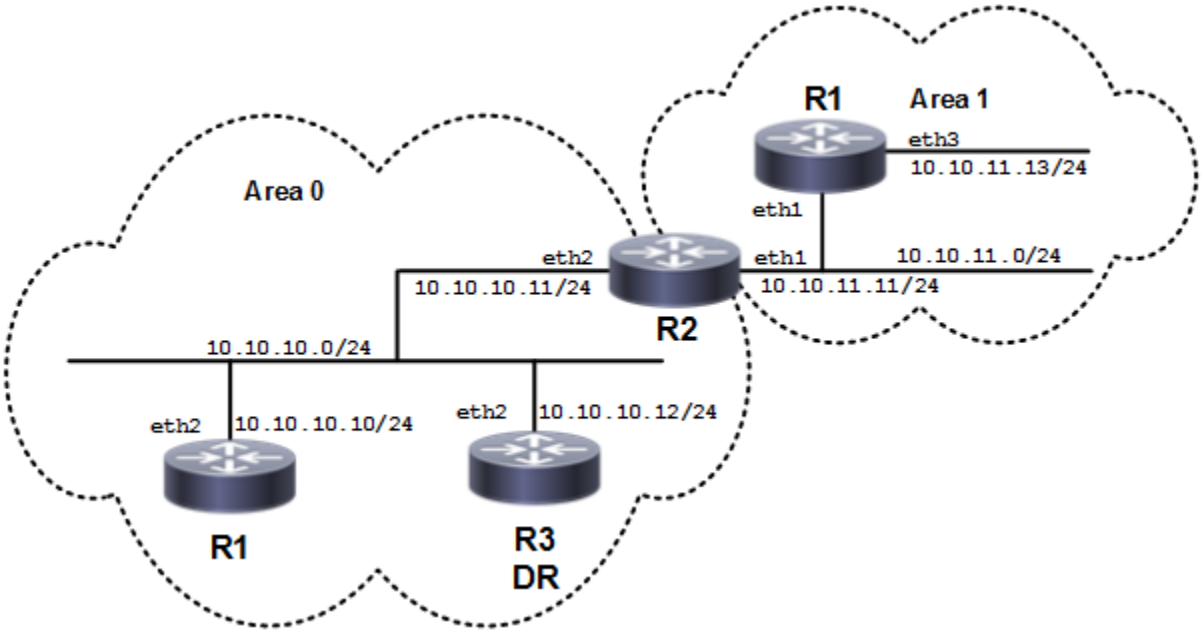

Figure 1-30: OSPF ABR Topology

Configuration

R2

`#configure terminal`	Enter configure mode
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer.
`(config-router)#network 10.10.10.0/24 area 0`	Define one interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface.
`(config-router)#network 10.10.11.0/24 area 1`	Define the other interface `(10.10.11.0/24)` on which OSPF runs, and associate the area ID `(1)` with the interface.

Validation

R2

```
#show running-config
!
no service password-encryption
!
hostname R2
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.89/24
!
interface eth1
 ip address 10.10.11.11/24
!
interface eth2
 ip address 10.10.10.11/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
```

```
router ospf 100
 network 10.10.10.0/24 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.1
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

#sh ip ospf
 Routing Process "ospf 100" with ID 10.12.26.89
 Process uptime is 4 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incomming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 9
 Number of LSA received 18
 Number of areas attached to this router: 2
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 2
        Area has no authentication
        SPF algorithm last executed 00:01:54.085 ago
        SPF algorithm executed 7 times
        Number of LSA 11. Checksum 0x0428ac
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Number of fully adjacent virtual neighbors through this area is 0
        Area has no authentication
        SPF algorithm last executed 00:00:41.737 ago
SPF algorithm executed 3 times
        Number of LSA 8. Checksum 0x043ce4
 Dste Staus: Disabled

#show ip ospf interface
eth2 is up, line protocol is up
```

```
  Internet Address 10.10.10.11/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.89, Interface Address 10.10.10.11
  Backup Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:11
  Neighbor Count is 2, Adjacent neighbor count is 2
  Suppress hello for 0 neighbor(s)
  Hello received 66 sent 38, DD received 11 sent 7
  LS-Req received 2 sent 2, LS-Upd received 15 sent 14
  LS-Ack received 14 sent 10, Discarded 0
  No authentication
eth1 is up, line protocol is up
  Internet Address 10.10.11.11/24, Area 0.0.0.1, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.92, Interface Address 10.10.11.13
  Backup Designated Router (ID) 10.12.26.89, Interface Address 10.10.11.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:03
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 22 sent 24, DD received 3 sent 9
  LS-Req received 1 sent 1, LS-Upd received 4 sent 5
  LS-Ack received 4 sent 3, Discarded 0
  No authentication

#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/Backup     00:00:34    10.10.10.10     eth2
0
10.12.26.90       1   Full/DROther    00:00:32    10.10.10.12     eth2
0
10.12.26.92       1   Full/DR         00:00:33    10.10.11.13     eth1
0

#show ip ospf database

            OSPF Router with ID (10.12.26.89) (Process ID 100 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router      Age     Seq#       CkSum  Link count
10.12.26.88     10.12.26.88     365     0x80000005 0x10bc 1
10.12.26.89     10.12.26.89     312     0x80000006 0x0fb8 1
10.12.26.90     10.12.26.90     363     0x80000003 0x10b8 1

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router      Age     Seq#       CkSum
10.10.10.11     10.12.26.89     364     0x80000002 0xe7fd
```

```
                Summary Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#     CkSum     Route
10.10.11.0      10.12.26.89      312     0x80000001 0x95fd 10.10.11.0/24

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#     CkSum  Opaque ID
1.0.0.1         10.12.26.88      363     0x80000003 0xa972 1
1.0.0.1         10.12.26.89      362     0x80000003 0xad6c 1
1.0.0.1         10.12.26.90      363     0x80000001 0xb564 1
1.0.0.10        10.12.26.88      363     0x80000003 0x0a32 10
1.0.0.10        10.12.26.89      362     0x80000002 0x2417 10
1.0.0.10        10.12.26.90      363     0x80000001 0x3efb 10

                Router Link States (Area 0.0.0.1)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
10.12.26.89     10.12.26.89      245     0x80000004 0x3d88 1
10.12.26.92     10.12.26.92      241     0x80000004 0x2698 1

                Net Link States (Area 0.0.0.1)
Link ID         ADV Router       Age     Seq#      CkSum
10.10.11.13     10.12.26.92      246     0x80000001 0x6ffb

                Summary Link States (Area 0.0.0.1)

Link ID         ADV Router       Age     Seq#      CkSum     Route
10.10.10.0      10.12.26.89      312     0x80000001 0xa0f3 10.10.10.0/24

                Area-Local Opaque-LSA (Area 0.0.0.1)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         10.12.26.89      243     0x80000001 0xb16a 1
1.0.0.1         10.12.26.92      244     0x80000001 0xbd58 1
1.0.0.8         10.12.26.89      234     0x80000002 0x96a2 8
1.0.0.8         10.12.26.92      244     0x80000001 0xc272 8
```

Redistribute Routes into OSPF

In this example, the configuration causes BGP routes to be imported into the OSPF routing table, and advertised as Type 5 External LSAs into Area 0.

Topology

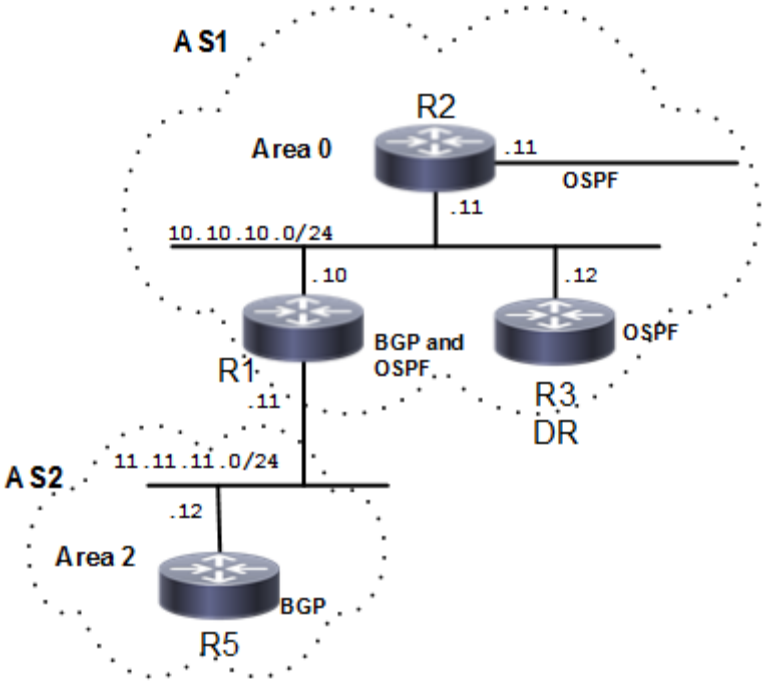

Figure 1-31: Redistribute Routes

R1

`#configure terminal`	Enter configure mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define one interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).
`(config-router)#redistribute bgp`	Specify redistributing routes from other routing protocol (BGP) into OSPF.

Validation

```
#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

C  10.10.9.0/24 [1] is directly connected, eth2, Area 0.0.0.0
C  10.10.10.0/24 [1] is directly connected, eth3, Area 0.0.0.0
O  10.10.11.0/24 [101] via 10.10.10.11, eth3, Area 0.0.0.0
C  10.10.12.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  10.10.13.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0
O  10.10.14.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0
```

Cost

A route can be made the preferred route by changing its cost. In this example, cost has been configured to make R2 the next hop for R1.

The default cost for each interface is 1. Interface eth2 on R2 has a cost of 100, and Interface eth2 on R3 has a cost of 150. The total cost to reach 10.10.14.0/24 (R4) through R2 and R3 is computed as follows:

R2: 1+100 = 101

R3: 1+150 = 151

Therefore, R1 chooses R2 as its next hop to destination 10.10.14.0/24 because it has the lower cost.

Topology

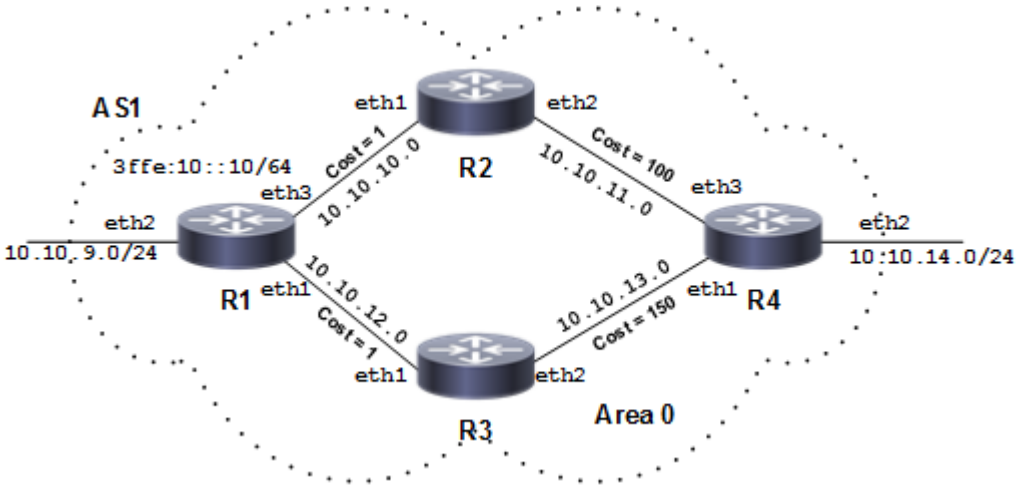

Figure 1-32: Configure Cost Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.

`(config-router)#network 10.10.9.0/24 area 0`
`(config-router)#network 10.10.10.0/24 area 0`
`(config-router)#network 10.10.12.0/24 area 0`
	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface (area ID `0` specifies the backbone area).

R2

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip ospf cost 100`	Set the OSPF cost of this link to `100`.
`(config-if)#exit`	Exit interface mode.

`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0` `(config-router)#network 10.10.11.0/24 area 0`	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface.

R3

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip ospf cost 150`	Set the OSPF cost of this link to `100`.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.12.0/24 area 0` `(config-router)#network 10.10.13.0/24 area 0`	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface.

R4

`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.

`(config-router)#network 10.10.11.0/24 area 0`
`(config-router)#network 10.10.13.0/24 area 0`
`(config-router)#network 10.10.14.0/24 area 0`
	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface.

Validation

R1

```
#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

C  10.10.9.0/24 [1] is directly connected, eth2, Area 0.0.0.0
C  10.10.10.0/24 [1] is directly connected, eth3, Area 0.0.0.0
O  10.10.11.0/24 [101] via 10.10.10.11, eth3, Area 0.0.0.0
C  10.10.12.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  10.10.13.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0
O  10.10.14.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0

#sh ip ospf interface
eth3 is up, line protocol is up
  Internet Address 10.10.10.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.88, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
```

```
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Backup Designated Router (ID) 10.12.26.89, Interface Address 10.10.10.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:03
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 43 sent 69, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 16 sent 18
  LS-Ack received 10 sent 11, Discarded 0
  No authentication
eth2 is up, line protocol is up
  Internet Address 10.10.9.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.88, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.9.10
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
  Hello received 0 sent 68, DD received 0 sent 0
  LS-Req received 0 sent 0, LS-Upd received 0 sent 0
  LS-Ack received 0 sent 0, Discarded 0
  No authentication
eth1 is up, line protocol is up
  Internet Address 10.10.12.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.88, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.12.10
  Backup Designated Router (ID) 10.12.26.90, Interface Address 10.10.12.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
Hello received 29 sent 66, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 10 sent 12
  LS-Ack received 10 sent 9, Discarded 0
  No authentication
```

R2

```
#sh ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

O  10.10.9.0/24 [2] via 10.10.10.10, eth1, Area 0.0.0.0
C  10.10.10.0/24 [1] is directly connected, eth1, Area 0.0.0.0
C  10.10.11.0/24 [100] is directly connected, eth2, Area 0.0.0.0
O  10.10.12.0/24 [2] via 10.10.10.10, eth1, Area 0.0.0.0
O  10.10.13.0/24 [101] via 10.10.11.11, eth2, Area 0.0.0.0
O  10.10.14.0/24 [101] via 10.10.11.11, eth2, Area 0.0.0.0

#sh ip ospf interface
```

```
eth2 is up, line protocol is up
  Internet Address 10.10.11.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 100
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 100
  Designated Router (ID) 10.12.26.89, Interface Address 10.10.11.10
  Backup Designated Router (ID) 10.12.26.92, Interface Address 10.10.11.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 56 sent 77, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 11 sent 7
  LS-Ack received 4 sent 8, Discarded 0
  No authentication
eth1 is up, line protocol is up
  Internet Address 10.10.10.11/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.89, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.10.10
  Backup Designated Router (ID) 10.12.26.89, Interface Address 10.10.10.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 74 sent 75, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 18 sent 16
  LS-Ack received 10 sent 12, Discarded 0
  No authentication
```

R3

```
#sh ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

O  10.10.9.0/24 [2] via 10.10.12.10, eth1, Area 0.0.0.0
O  10.10.10.0/24 [2] via 10.10.12.10, eth1, Area 0.0.0.0
O  10.10.11.0/24 [102] via 10.10.12.10, eth1, Area 0.0.0.0
C  10.10.12.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  10.10.13.0/24 [103] via 10.10.12.10, eth1, Area 0.0.0.0
O  10.10.14.0/24 [103] via 10.10.12.10, eth1, Area 0.0.0.0

#sh ip ospf interface
eth2 is up, line protocol is up
  Internet Address 10.10.13.10/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.90, Network Type
BROADCAST, Cost: 150
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 150
  Designated Router (ID) 10.12.26.90, Interface Address 10.10.13.10
  Backup Designated Router (ID) 10.12.26.92, Interface Address 10.10.13.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
```

© 2025 IP Infusion Inc. Proprietary

```
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 85 sent 94, DD received 3 sent 4
  LS-Req received 0 sent 0, LS-Upd received 3 sent 4
  LS-Ack received 3 sent 3, Discarded 0
  No authentication
eth1 is up, line protocol is up
  Internet Address 10.10.12.11/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.90, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 10.10.12.10
  Backup Designated Router (ID) 10.12.26.90, Interface Address 10.10.12.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 92 sent 92, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 12 sent 10
  LS-Ack received 8 sent 10, Discarded 0
  No authentication
```

R4

```
#sh ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

O  10.10.9.0/24 [3] via 10.10.11.10, eth3, Area 0.0.0.0
                    via 10.10.13.10, eth1, Area 0.0.0.0
O  10.10.10.0/24 [2] via 10.10.11.10, eth3, Area 0.0.0.0
C  10.10.11.0/24 [1] is directly connected, eth3, Area 0.0.0.0
O  10.10.12.0/24 [2] via 10.10.13.10, eth1, Area 0.0.0.0
C  10.10.13.0/24 [1] is directly connected, eth1, Area 0.0.0.0
C  10.10.14.0/24 [1] is directly connected, eth2, Area 0.0.0.0

#sh ip ospf interface
eth3 is up, line protocol is up
  Internet Address 10.10.11.11/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.92, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.89, Interface Address 10.10.11.10
  Backup Designated Router (ID) 10.12.26.92, Interface Address 10.10.11.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 95 sent 96, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 7 sent 11
  LS-Ack received 7 sent 5, Discarded 0
  No authentication
eth2 is up, line protocol is up
  Internet Address 10.10.14.10/24, Area 0.0.0.0, MTU 1500
```

```
  Process ID 100, VRF (default), Router ID 10.12.26.92, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.92, Interface Address 10.10.14.10
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:10
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
  Hello received 0 sent 95, DD received 0 sent 0
  LS-Req received 0 sent 0, LS-Upd received 0 sent 0
  LS-Ack received 0 sent 0, Discarded 0
  No authentication
eth1 is up, line protocol is up
  Internet Address 10.10.13.11/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 10.12.26.92, Network Type
BROADCAST, Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.90, Interface Address 10.10.13.10
  Backup Designated Router (ID) 10.12.26.92, Interface Address 10.10.13.11
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
 Hello received 92 sent 93, DD received 4 sent 3
  LS-Req received 0 sent 0, LS-Upd received 4 sent 3
  LS-Ack received 3 sent 3, Discarded 0
  No authentication
```

Path MTU for IPv4 and IPv6

In IPv4, path MTU discovery enables a host to actively identify and adapt to variations in the MTU size across different links along a data path. In contrast, IPv6 adopts an approach where fragmentation is managed by the packet's source when the path MTU of a specific link along the data path cannot accommodate the packet's size. This approach, where IPv6 hosts handle packet fragmentation, conserves processing resources in IPv6 devices and enhances the overall efficiency of IPv6 networks.

Topology

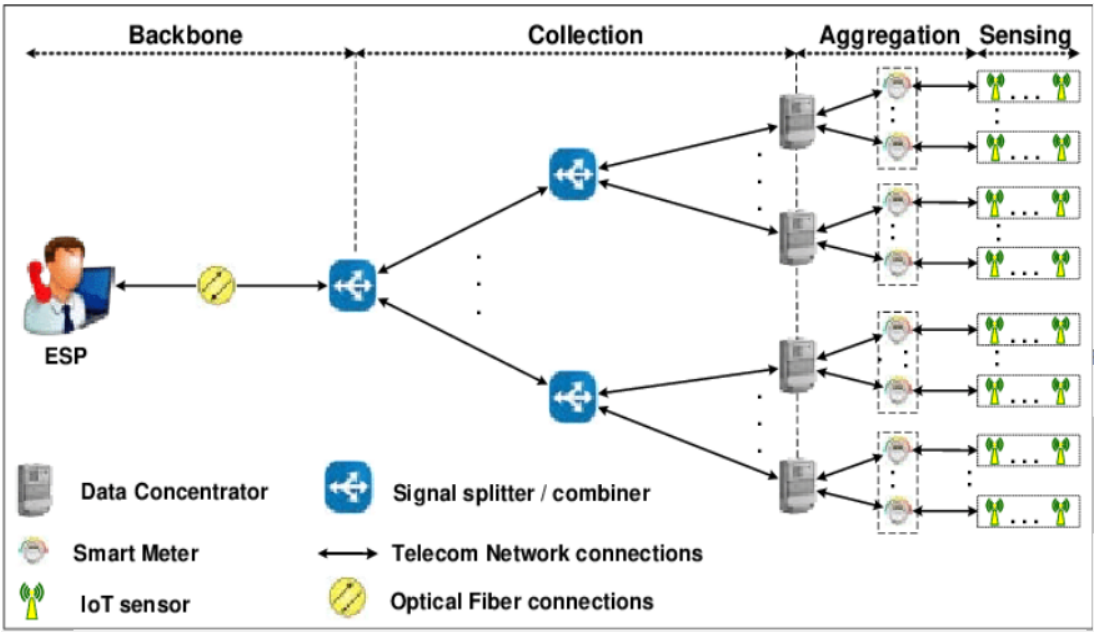

Figure 1-33: Path MTU Topology

R1

`OcNOS#configure terminal`	Enter configuration mode.
`OcNOS(config)#interface cd31/1`	Entering in to interface
`OcNOS(config-if)#mtu 1500`	Configure mtu
`OcNOS(config-if)#ip address 12.12.12.1/24`	Configure ipv4 address
`OcNOS(config-if)#ipv6 address 1001::1/64`	Configure ipv6 address
`OcNOS(config)#router ospf 1`	Configure ip ospf
`OcNOS(config-router)#ospf router-id 20.20.20.1`	Configure router id under ospf
`OcNOS(config-router)#network 12.12.12.0/24 area 0.0.0.0`	Add network under ospf
`OcNOS(config)#router ipv6 ospf 1`	Configure ipv6 ospf
`OcNOS(config-router)#router-id 1.1.1.1`	Configure router id under ospf
`OcNOS(config-router)#commit`	Commit all the transactions
`OcNOS(config)#exit`	Exit

R2

`OcNOS#configure terminal`	Enter configuration mode.
`OcNOS(config)#interface cd31/1`	Entering in to interface
`OcNOS(config-if)#mtu 1500`	Configure mtu
`OcNOS(config-if)#ip address 12.12.12.1/24`	Configure ipv4 address
`OcNOS(config-if)#ipv6 address 1001::1/64`	Configure ipv6 address

© 2025 IP Infusion Inc. Proprietary

OcNOS(config)#router ospf 1	Configure ip ospf
OcNOS(config-router)#ospf router-id 20.20.20.1	Configure router id under ospf
OcNOS(config-router)#network 12.12.12.0/24 area 0.0.0.0	Add network under ospf
OcNOS(config)#router ipv6 ospf 1	Configure ipv6 ospf
OcNOS(config-router)#router-id 1.1.1.1	Configure router id under ospf
OcNOS(config-router)#commit	Commit all the transactions
OcNOS(config)#exit	Exit

Validation

```
OcNOS#sh in
installers  interface
OcNOS#sh interface cd31/1
Interface cd31/1
  Flexport: Non Control Port (Active)
  Hardware is ETH  Current HW addr: e49d.7356.df24
  Physical:e49d.7356.df24  Logical:(not set)
  Forward Error Correction (FEC) configured is Auto (default)
  FEC status is N/A
  Port Mode is Router
  Protected Mode is Promiscuous
  Interface index: 10142
  Metric 1 mtu 1700 duplex-full  link-speed 10g
  Debounce timer: disable
  ARP ageing timeout 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: 2023 Apr 03 17:24:58 (18:36:43 ago)
  inet6 fe80::e69d:73ff:fe56:df24/64
  ND router advertisements are sent approximately every 411 seconds
  ND next router advertisement due in 176 seconds.
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
  5 minute input rate 3 bits/sec, 0 packets/sec
  5 minute output rate 1 bits/sec, 0 packets/sec
  RX
    unicast packets 0 multicast packets 180 broadcast packets 0
    input packets 180 bytes 15000
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
  TX
    unicast packets 0 multicast packets 176 broadcast packets 0
    output packets 176 bytes 14632
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0
```

Virtual Links

Virtual links are used to connect a temporarily-disjointed non-backbone area to the backbone area, or to repair a non-contiguous backbone area. In this example, the ABR R3 has temporarily lost connection to Area 0, in turn, disconnecting Area 2 from the backbone area. The virtual link between ABR R1 and ABR R2 connects Area 2 to Area 0. Area 1 is used as a transit area.

Topology

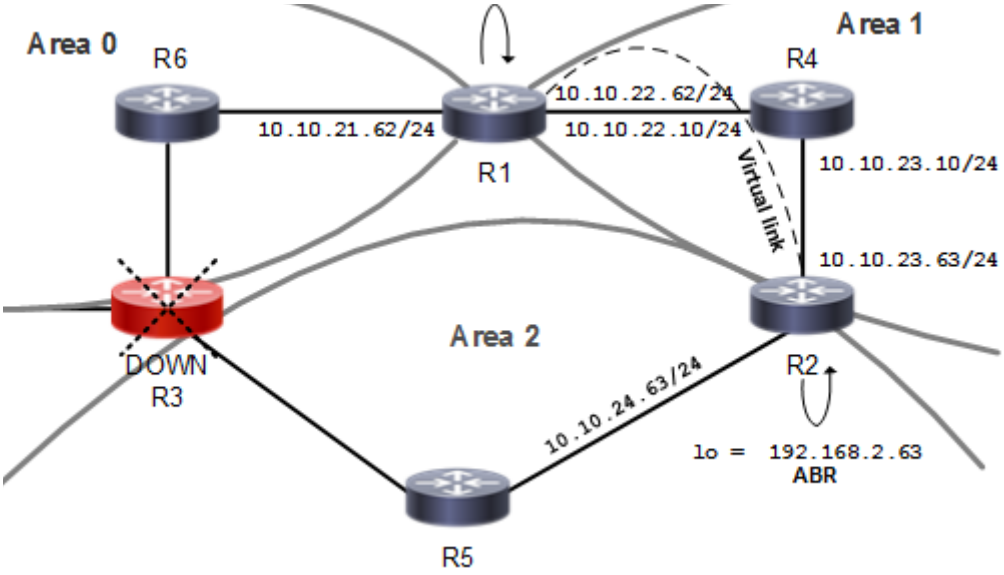

Figure 1-34: Virtual Links Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify loopback as the interface you want to configure.
`(config-if)#ip address 192.168.1.62/32`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#ospf router-id 192.168.1.62`	Configure the OSPF Router ID (`192.168.1.62`) for this router.
`(config-router)#network 10.10.21.0/24 area 0` `(config-router)#network 10.10.22.0/24 area 1`	Define interfaces on which OSPF runs, and associate the area IDs `(0 and 1)` with the interface.
`(config-router)#area 1 virtual-link 192.168.2.63`	Configure a virtual link between this router R1 and R2 (Router ID `192.168.2.63`) through transit area 1.

R2

`(config)#interface lo`	Specify loopback as the interface you want to configure.
`(config-if)#ip address 192.168.2.63/32`	Configure the IP address on this interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#ospf router-id 192.168.2.63`	Configure the OSPF Router ID (`192.168.1.63`) for this router.

`(config-router)#network 10.10.23.0/24 area 1`
`(config-router)#network 10.10.24.0/24 area 2`
`(config-router)#network 192.168.2.63/32 area 2`
	Define interfaces on which OSPF runs, and associate the area IDs `(1 and 2)` with the interface.

`(config-router)#area 1 virtual-link 192.168.1.62`	Configure a virtual link between this router R2 and R1 (Router ID `192.168.2.62`) through transit area `1`.

Validation

```
R1#show ip ospf virtual-links
Virtual Link VLINK0 to router 2.2.2.2 is up
  Transit area 0.0.0.1 via interface eth2
  Hello suppression enabled
  DoNotAge LSA allowed
  Local address 13.13.13.1/32
  Remote address 12.12.12.1/32
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  No authentication
    Adjacency state Down

R2#show ip ospf virtual-links
Virtual Link VLINK0 to router 1.1.1.1 is up
  Transit area 0.0.0.1 via interface eth1
  Hello suppression enabled
  DoNotAge LSA allowed
  Local address 12.12.12.1/32
  Remote address 13.13.13.1/32
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  No authentication
    Adjacency state Init

R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
192.168.20.5      1   Full/DR         00:00:34    13.13.13.2      eth2
0
```

```
R2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
192.168.20.5      1   Full/DR         00:00:36    12.12.12.2      eth1
0
1.1.1.1           1   Init/ -           00:00:32    13.13.13.1      VLINK0
R1#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

IA 2.2.2.2/32 [12] via 13.13.13.2, eth2, Area 0.0.0.1
O  12.12.12.0/24 [2] via 13.13.13.2, eth2, Area 0.0.0.1
C  13.13.13.0/24 [1] is directly connected, eth2, Area 0.0.0.1

R2#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

C  2.2.2.2/32 [10] is directly connected, lo, Area 0.0.0.2
C  12.12.12.0/24 [1] is directly connected, eth1, Area 0.0.0.1
O  13.13.13.0/24 [2] via 12.12.12.2, eth1, Area 0.0.0.1

R1#show ip ospf
 Routing Process "ospf 100" with ID 1.1.1.1
 Process uptime is 39 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 6
```

© 2025 IP Infusion Inc. Proprietary

```
 Number of LSA received 15
 Number of areas attached to this router: 2
MemPool - struct ospf lsa             : (0-16) | Total (16/100000)
blk_size:160
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 0
        Area has no authentication
        SPF algorithm last executed 00:10:05.434 ago
        SPF algorithm executed 1 times
        Number of LSA 3. Checksum 0x01bf9c
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Number of fully adjacent virtual neighbors through this area is 0
        Area has no authentication
        SPF algorithm last executed 00:09:57.432 ago
        SPF algorithm executed 7 times
        Number of LSA 13. Checksum 0x076e78
 Dste Staus: Disabled

R2#show ip ospf
 Routing Process "ospf 100" with ID 2.2.2.2
 Process uptime is 16 hours 48 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 11
 Number of LSA received 12
 Number of areas attached to this router: 3
MemPool - struct ospf lsa             : (0-20) | Total (20/100000)
blk_size:160
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 0
        Area has no authentication
        SPF algorithm last executed 00:11:05.618 ago
        SPF algorithm executed 1 times
```

```
        Number of LSA 4. Checksum 0x018ce2
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Number of fully adjacent virtual neighbors through this area is 0
        Area has no authentication
        SPF algorithm last executed 00:11:03.619 ago
        SPF algorithm executed 6 times
        Number of LSA 13. Checksum 0x076e78
    Area 0.0.0.2
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 0
        Number of fully adjacent virtual neighbors through this area is 0
        Area has no authentication
        SPF algorithm last executed 00:11:05.618 ago
        SPF algorithm executed 3 times
        Number of LSA 3. Checksum 0x0139cf
Dste Staus: Disabled
```

OSPF Authentication

There are three types of OSPF authentications--Null (Type 0), Simple Text (Type 1), and MD5 (Type 2). With Null authentication, routing exchanges over the network are not authenticated. In Simple Text authentication, the authentication type is the same for all routers that communicate using OSPF in a network. For MD5 authentication, configure a key and a key ID on each router. The router generates a message digest on the basis of the key, key ID, and OSPF packet, and adds it to the OSPF packet.

The authentication type can be configured on a per-interface basis or a per-area basis. Additionally, Interface and Area authentication can be used together. Area authentication is used for an area, and interface authentication is used for a specific interface in the area. If the Interface authentication type is different from the Area authentication type, the Interface authentication type overrides the Area authentication type. If the Authentication type is not specified for an interface, the Authentication type for the area is used. The authentication command descriptions contain details of each type of authentication.

In the example below, R1 and R2 are configured for both the interface and area authentications. The authentication type of interface eth1 on R1 and interface eth0 on R2 is MD5 mode, and is defined by the area authentication command; however, the authentication type of interface eth2 on R1 and interface eth1 on R2 is plain text mode, and is defined by the `ip ospf authentication` command. This interface command overrides the `area authentication` command.

Topology

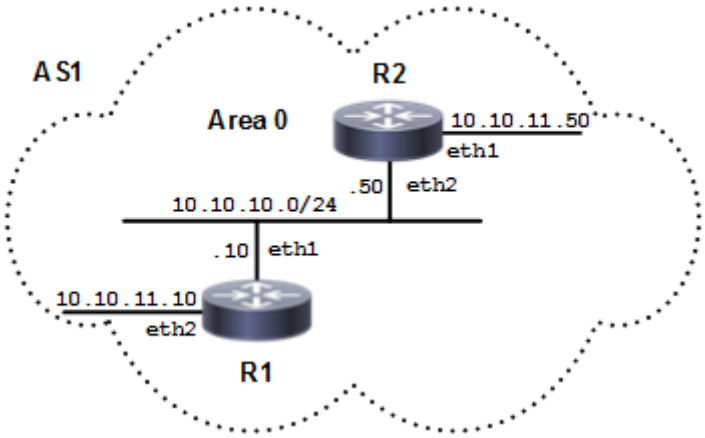

Figure 1-35: OSPF Authentication Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0` `(config-router)#network 10.10.11.0/24 area 0`	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface (area ID 0 specifies the backbone area).
`(config-router)#area 0 authentication message-digest`	Enable MD5 authentication on area 0.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip ospf message-digest-key 1 md5 test`	Register the MD5 key test for OSPF authentication. The key ID is 1.
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip ospf authentication`	Enable the OSPF packet to use text authentication on the current interface (`eth2`).
`(config-if)#ip ospf authentication-key test`	Specify an OSPF authentication password (`test`) for the neighboring routers.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 100`	Configure the routing process, and specify the Process ID (`100`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0` `(config-router)#network 10.10.11.0/24 area 0`	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface (area ID 0 specifies the backbone area).

© 2025 IP Infusion Inc. Proprietary

`(config-router)#area 0 authentication message-digest`	Enable MD5 authentication on area 0.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip ospf message-digest-key 1 md5 test`	Register MD5 key test for OSPF authentication. The key ID is 1.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip ospf authentication`	Enable the OSPF packet to use text authentication on the current interface (`eth1`).
`(config-if)#ip ospf authentication-key test`	Specify an OSPF authentication password `test` for the neighboring routers.

Validation

R1

```
R1#sh running-config
!
no service password-encryption
!
hostname R1
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.88/24
!
interface eth1
 ip address 10.10.10.10/24
 ip ospf message-digest-key 1 md5 0x293da85becc67703
!
interface eth2
 ip address 10.10.11.10/24
 ip ospf authentication
 ip ospf authentication-key 0x293da85becc67703
!
interface eth3
!
interface eth4
!
interface eth5
```

```
!
interface eth6
!
interface eth7
!
router ospf 100
 area 0.0.0.0 authentication message-digest
 network 10.10.9.0/24 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.0
 network 10.10.12.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end


R1#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.89       1   Full/DR         00:00:38    10.10.10.50     eth1
0
```

R2

```
R2#sh running-config
!
no service password-encryption
!
hostname R2
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.89/24
!
interface eth1
 ip address 10.10.11.50/24
 ip ospf authentication
 ip ospf authentication-key 0x293da85becc67703
```

```
!
interface eth2
 ip address 10.10.10.50/24
 ip ospf message-digest-key 1 md5 0x293da85becc67703
 ip ospf cost 100
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
router ospf 100
 area 0.0.0.0 authentication message-digest
 network 10.10.10.0/24 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

R2#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/Backup     00:00:33    10.10.10.10     eth2
0
```

Multiple OSPF Instances

By using multiple OSPF instances, OSPF routes can be segregated, based on their instance number. Routes of one instance are stored differently from routes of another instance running in the same router.

To configure multiple OSPF instances, perform the following procedures referring to the topology diagram below:

1. Enable OSPF on an interface.
2. Enable multiple instances.
3. Configure redistribution among multiple instances.

   Note: Optionally, redistribution can be configured with the metric, type or route-map options.

Topology

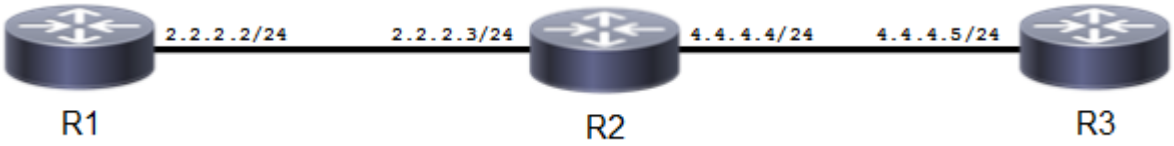

Figure 1-36: Multiple OSPF Instances

Enable Multiple OSPF Instances on a Router

In this example, routers R1, R2, and R3 are in Area 0, and all run OSPF.

R1

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ip address 2.2.2.2/24`	Specify the IP address of the interface.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 10`	Configure an OSPF instance with an instance ID of `10`.
`(config-router)#router-id 5.5.5.5`	Configure the router ID to use on this instance.
`(config-router)#network 2.2.2.0/24 area 0`	Advertise the network with the area ID.

R2

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ip address 2.2.2.3/24`	Specify the IP address of the interface.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 10`	Configure an OSPF instance with an instance ID of `10`.
`(config-router)#router-id 6.6.6.6`	Configure the router ID to use on this instance.
`(config-router)#network 2.2.2.0/24 area 0`	Advertise the network with the area ID.
`(config-router)#exit`	Exit router mode.
`(config)#interface eth2`	Enter interface mode for `eth2`.
`(config-if)#ip address 4.4.4.4/24`	Configure the IP address.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 15`	Configure an OSPF instance with an instance ID of `15`.
`(config-router)#router-id 8.8.8.8`	Configure the router ID to use on this instance.
`(config-router)#network 4.4.4.0/24 area 0`	Advertise the network with the area ID.

R3

`(config)#``interface eth1`	Enter interface mode for `eth1`.
`(config-if)#``ip address 4.4.4.5/24`	Configure the IP address.
`(config-if)#``no shutdown`	Activate the interface.
`(config-if)#``exit`	Exit interface mode.
`(config)#``router ospf 15`	Configure an OSPF instance with an instance ID of `15`.
`(config-router)#``router-id 7.7.7.7`	Configure the router ID to use on this instance.
`(config-router)#``network 4.4.4.0/24 area 0`	Advertise the network with the area ID.

© 2025 IP Infusion Inc. Proprietary

Validation

R1

```
R1#sh running-config
!
no service password-encryption
!
hostname R1
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.88/24
!
interface eth1
 ip address 2.2.2.2/24
!
interface eth2
 ip address 10.10.11.10/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
interface eth7
!
router ospf 10
 ospf router-id 5.5.5.5
 network 2.2.2.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

R1#sh ip ospf neighbor
```

```
Total number of full neighbors: 1
OSPF process 10 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
6.6.6.6           1   Full/Backup     00:00:39    2.2.2.3         eth1
0

R1#sh ip ospf route

OSPF process 10:
Codes: C - connected, D - Discard, 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

C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0
```

R2

```
R2#sh running-config
!
no service password-encryption
!
hostname R2
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.89/24
!
interface eth1
 ip address 2.2.2.3/24
!
interface eth2
 ip address 4.4.4.4/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
router ospf 10
 ospf router-id 6.6.6.6
 network 2.2.2.0/24 area 0.0.0.0
cspf disable-better-protection
```

```
!
router ospf 15
 ospf router-id 8.8.8.8
 network 4.4.4.0/24 area 0.0.0.0
 no capability cspf
!
line con 0
 login
line vty 0 39
 login
!
end

R2#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 10 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
5.5.5.5           1   Full/DR         00:00:33    2.2.2.2         eth1
0

Total number of full neighbors: 1
OSPF process 15 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
7.7.7.7           1   Full/Backup     00:00:31    4.4.4.5         eth2
0

R2#sh ip ospf route

OSPF process 10:
Codes: C - connected, D - Discard, 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

C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0

OSPF process 15:
Codes: C - connected, D - Discard, 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

C  4.4.4.0/24 [1] is directly connected, eth2, Area 0.0.0.0
```

R3

```
R3#sh running-config
!
no service password-encryption
!
hostname R3
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
```

```
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.26.90/24
!
interface eth1
 ip address 4.4.4.5/24
!
interface eth2
 ip address 10.10.13.10/24
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
router ospf 15
 ospf router-id 7.7.7.7
 network 4.4.4.0/24 area 0.0.0.0
 cspf disable-better-protection
!
line con 0
 login
line vty 0 39
 login
!
end

R3#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 15 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
8.8.8.8           1   Full/DR         00:00:30    4.4.4.4         eth1
0
```

Redistribute among Multiple Instances

In this example, routes of one instance are redistributed to another instance to enable ping from R1 to R3 or vice versa; and R2 redistributes routes from one instance to another.

R2

`(config)#router ospf 15`	Configure an OSPF instance with instance ID `15`.
`(config-router)#router-id 8.8.8.8`	Configure the router ID.

(config-router)#redistribute ospf 10	Redistribute instance 10 routes.
(config-router)#redistribute connected	Redistribute connected routes to instance 15.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 10	Configure an OSPF instance with instance ID 10.
(config-router)#router-id 6.6.6.6	Configure the router ID.
(config-router)#redistribute ospf 15	Redistribute instance 15 routes.
(config-router)#redistribute connected	Redistribute connected routes to instance 10.

Redistribute with the Metric Option

In this example, on R3, R1 and R2 have each other's routes with a metric of 100.

R2

(config)#router ospf 15	Configure an OSPF instance with instance ID 15.
(config-router)#router-id 8.8.8.8	Configure the router ID.
(config-router)#redistribute ospf 10 metric 100	Redistribute instance 10 routes with metric 100.
(config-router)#redistribute connected	Redistribute connected routes to instance 15.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 10	Configure an OSPF instance with instance ID 10.
(config-router)#router-id 6.6.6.6	Configure the router ID.
(config-router)#redistribute ospf 15 metric 100	Redistribute instance 15 routes with metric 100.
(config-router)#redistribute connected	Redistribute connected routes to instance 10.

Redistribute with the Type Option

In this example, on R3, R1 has R3 routes as type 2, and R2 has R1 routes as type 1.

R2

(config)#router ospf 15	Configure an OSPF instance with instance ID 15.
(config-router)#router-id 8.8.8.8	Configure the router ID.
(config-router)#redistribute ospf 10 metric-type 1	Redistribute instance 10 routes with metric-type 1.
(config-router)#redistribute connected	Redistribute connected routes to instance 15.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 10	Configure an OSPF instance with instance ID 10.
(config-router)#router-id 6.6.6.6	Configure the router ID.
(config-router)#redistribute ospf 15 metric-type 2	Redistribute instance 15 routes with type 2.
(config-router)#redistribute connected	Redistribute connected routes to instance 10.

Redistribute with the Route-Map Option

R2

`(config)#route-map 1 permit 10`	Enter route-map mode, specifying route-map ID.
`(config-route-map)#set metric 100`	Set metric value.
`(config-route-map)#set metric-type type-2`	Set metric-type.
`(config-route-map)#exit`	Exit route-map mode.
`(config)#route-map 2 permit 10`	Enter route-map mode, specifying route-map ID.
`(config-route-map)#set metric 200`	Set metric value.
`(config-route-map)#set metric-type type-1`	Set metric-type.
`(config-route-map)#exit`	Exit route-map mode.
`(config)#router ospf 15`	Configure an OSPF instance with instance ID `15`.
`(config-router)#router-id 8.8.8.8`	Configure the router ID.
`(config-router)#redistribute ospf 10 route-map 1`	Redistribute instance `10` routes with route map `1`.
`(config-router)#redistribute connected`	Redistribute connected routes to instance `15`.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.
`(config)#router ospf 10`	Configure an OSPF instance with instance ID `10`.
`(config-router)#router-id 6.6.6.6`	Configure the router ID.
`(config-router)#redistribute ospf 15 route-map 2`	Redistribute instance `15` routes with route map `2`.
`(config-router)#redistribute connected`	Redistribute connected routes to instance `10`.

Validation

```
R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 10 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
6.6.6.6           1   Full/DR         00:00:39    2.2.2.3       eth1
0

R2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 10 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
5.5.5.5           1   Full/Backup     00:00:35    2.2.2.2       eth1
0

Total number of full neighbors: 1
OSPF process 15 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
```

```
7.7.7.7         1   Full/Backup     00:00:36    4.4.4.5        eth2
0
R3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 15 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
8.8.8.8         1   Full/DR         00:00:40    4.4.4.4        eth2
0
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     2.2.2.0/24 is directly connected, eth1, 00:08:40
O E1  4.4.4.0/24 [110/201] via 2.2.2.3, eth1, 00:01:18
C     5.5.5.5/32 is directly connected, lo, 00:08:41
O E2  6.6.6.6/32 [110/20] via 2.2.2.3, eth1, 00:01:10
O E2  8.8.8.8/32 [110/20] via 2.2.2.3, eth1, 00:01:10
C     127.0.0.0/8 is directly connected, lo, 00:08:44
C     192.168.20.0/24 is directly connected, eth0, 00:08:40

Gateway of last resort is not set

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     2.2.2.0/24 is directly connected, eth1, 5d00h02m
C     4.4.4.0/24 is directly connected, eth2, 5d00h02m
C     6.6.6.6/32 is directly connected, lo, 4d23h59m
C     8.8.8.8/32 is directly connected, lo, 4d23h59m
C     127.0.0.0/8 is directly connected, lo, 5d00h09m
C     192.168.20.0/24 is directly connected, eth0, 5d00h08m

Gateway of last resort is not set

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O E2  2.2.2.0/24 [110/20] via 4.4.4.4, eth2, 00:02:45
C     4.4.4.0/24 is directly connected, eth2, 00:07:12
C     5.5.5.5/32 is directly connected, lo, 00:16:35
O E2  6.6.6.6/32 [110/20] via 4.4.4.4, eth2, 00:02:45
O E2  8.8.8.8/32 [110/20] via 4.4.4.4, eth2, 00:02:45
C     127.0.0.0/8 is directly connected, lo, 00:16:39
C     192.168.20.0/24 is directly connected, eth0, 00:15:36

Gateway of last resort is not set

#show ip ospf route

OSPF process 100:
Codes: C - connected, D - Discard, 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

C  10.10.9.0/24 [1] is directly connected, eth2, Area 0.0.0.0
C  10.10.10.0/24 [1] is directly connected, eth3, Area 0.0.0.0
O  10.10.11.0/24 [101] via 10.10.10.11, eth3, Area 0.0.0.0
C  10.10.12.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  10.10.13.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0
O  10.10.14.0/24 [102] via 10.10.10.11, eth3, Area 0.0.0.0

R2#show route-map

route-map 1, permit, sequence 10
  Match clauses:
  Set clauses:
    metric 100
    metric-type type-2
route-map 2, permit, sequence 10
  Match clauses:
  Set clauses:
    metric 200
    metric-type type-1

R1#show ip ospf route

OSPF process 10:
Codes: C - connected, D - Discard, 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

C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0
E1 4.4.4.0/24 [201] via 2.2.2.3, eth1
E2 6.6.6.6/32 [1/20] via 2.2.2.3, eth1
E2 8.8.8.8/32 [1/20] via 2.2.2.3, eth1
E2 192.168.20.0/24 [1/20] via 2.2.2.3, eth1

R2#show ip ospf route
```

```
OSPF process 10:
Codes: C - connected, D - Discard, 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

C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0

OSPF process 15:
Codes: C - connected, D - Discard, 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

C  4.4.4.0/24 [1] is directly connected, eth2, Area 0.0.0.0

R3#show ip ospf route

OSPF process 15:
Codes: C - connected, D - Discard, 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

E2 2.2.2.0/24 [1/20] via 4.4.4.4, eth2
C  4.4.4.0/24 [1] is directly connected, eth2, Area 0.0.0.0
E2 6.6.6.6/32 [1/20] via 4.4.4.4, eth2
E2 8.8.8.8/32 [1/20] via 4.4.4.4, eth2
E2 192.168.20.0/24 [1/20] via 4.4.4.4, eth2
```

Multiple OSPF Instances on Same Subnet

Multiple OSPF instances can be configured on the same subnet. The OSPF instance ID supports separate OSPFv2 protocol instances. With this feature, an adjacency is formed only if the received packet's instance ID is the same as the instance ID configured for that interface.

Topology

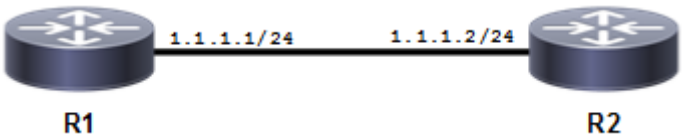

Figure 1-37: Multiple Instances on the Same Subnet

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#enable ext-ospf-multi-inst`	Enable multiple-instance capability.
`(config)#router ospf 1`	Configure an OSPF instance with an instance ID of `1`.

(config-router)#network 1.1.1.0/24 area 0 instance-id 1	Advertise the network in Area 0 with an instance ID of 1.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 2	Configure an OSPF instance with an instance ID of 2.
(config-router)#network 1.1.1.0/24 area 0 instance-id 2	Advertise the network in Area 0 with an instance ID of 2.
(config-router)#exit	Exit Router mode, and return to Configure mode.

R2

#configure terminal	Enter configure mode.
(config)#enable ext-ospf-multi-inst	Enable multiple-instance capability.
(config)#router ospf 1	Configure an OSPF instance with an instance ID of 1.
(config-router)#network 1.1.1.0/24 area 0 instance-id 1	Advertise the network in Area 0 with an instance ID of 1.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#router ospf 2	Configure an OSPF instance with an instance ID of 2.
(config-router)#network 1.1.1.0/24 area 0 instance-id 2	Advertise the network in Area 0 with an instance ID of 2.
(config-router)#exit	Exit Router mode, and return to Configure mode.

Validation

R1

```
R1#show ip ospf interface
eth1 is up, line protocol is up
  Internet Address 1.1.1.1/24, Area 0.0.0.0, MTU 1500
  Process ID 1, VRF (default), Router ID 10.12.26.88, Network Type BROADCAST,
Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 1.1.1.1
  Backup Designated Router (ID) 10.12.26.89, Interface Address 1.1.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:10
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 7 sent 16, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 3 sent 5
  LS-Ack received 3 sent 3, Discarded 0
  No authentication
  Internet Address 1.1.1.1/24, Area 0.0.0.0, MTU 1500
  Process ID 2, VRF (default), Router ID 10.12.26.88, Network Type BROADCAST,
Cost: 1
  Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 1.1.1.1
  Backup Designated Router (ID) 10.12.26.89, Interface Address 1.1.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
```

```
  Hello received 4 sent 12, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 3 sent 5
  LS-Ack received 3 sent 3, Discarded 0
  No authentication

R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.89       1   Full/Backup     00:00:35    1.1.1.2       eth1
1

Total number of full neighbors: 1
OSPF process 2 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.89       1   Full/Backup     00:00:33    1.1.1.2       eth1
2
```

R2

```
R2#sh ip ospf interface
eth1 is up, line protocol is up
  Internet Address 1.1.1.2/24, Area 0.0.0.0, MTU 1500
  Process ID 1, VRF (default), Router ID 10.12.26.89, Network Type BROADCAST,
Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 1.1.1.1
  Backup Designated Router (ID) 10.12.26.89, Interface Address 1.1.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 17 sent 17, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 5 sent 3
  LS-Ack received 2 sent 3, Discarded 0
  No authentication
  Internet Address 1.1.1.2/24, Area 0.0.0.0, MTU 1500
  Process ID 2, VRF (default), Router ID 10.12.26.89, Network Type BROADCAST,
Cost: 1
  Transmit Delay is 1 sec,  State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 10.12.26.88, Interface Address 1.1.1.1
  Backup Designated Router (ID) 10.12.26.89, Interface Address 1.1.1.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 13 sent 14, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 5 sent 3
  LS-Ack received 2 sent 3, Discarded 0
  No authentication

R2#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
```

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/DR         00:00:32    1.1.1.1         eth1
1

Total number of full neighbors: 1
OSPF process 2 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/DR         00:00:37    1.1.1.1         eth1
2
```

Multi-Area Adjacency Configuration

Multiple OSPF areas for a same subnet can be configured between two routers. In the diagram below, OSPF is enabled between R2 and R3 under area 0 and area 1, though there is only one link available between these two routers. Multi-area adjacency allows establishing adjacency on multiple areas between the Area Border Routers (ABRs). The specified interface of the ABR is associated with multiple areas.

Each multi-area-adjacency internally implements point-to-point functionality, once the adjacency reaches the FULL state. This point-to-point link provides a topological path for that area. Like a virtual link, there is no restriction for multi-area adjacency that the packets always go through the backbone.

Topology

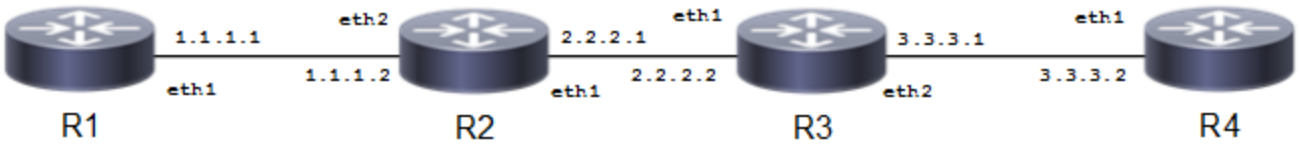

Figure 1-38: One Subnet with Multiple OSPF Areas

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure an OSPF instance with an instance ID of `1`.
`(config-router)#network 1.1.1.0/24 area 1`	Configure OSPF between R1 and R2 under area `1`.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure an OSPF instance with an instance ID of `1`.
`(config-router)#network 1.1.1.0/24 area 1`	Configure OSPF between R1 and R2 under area `1`.
`(config-router)#network 2.2.2.0/24 area 0`	Configure OSPF between R2 and R3 under area `0`.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.

(config)#interface eth1	Enter interface mode.
(config-if)#ip address 2.2.2.1/24	Configure IP address on the interface.
(config-if)#ip ospf 1 multi-area 0.0.0.1 neighbor 2.2.2.2	Configure multi area adjacency.
(config-if)#exit	Exit interface mode.

R3

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure an OSPF instance with an instance ID of 1.
(config-router)#network 2.2.2.0/24 area 0	Configure OSPF between R2 and R3 under area 0.
(config-router)#network 3.3.3.0/24 area 1	Configure OSPF between R3 and R4 under area 1.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 2.2.2.2/24	Configure IP address on the interface.
(config-if)#ip ospf 1 multi-area 0.0.0.1 neighbor 2.2.2.1	Configure multi area adjacency.
(config-if)#exit	Exit interface mode.

R4

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure an OSPF instance with an instance ID of 1.
(config-router)#network 3.3.3.0/24 area 1	Configure OSPF between R3 and R4 under area 1.
(config-router)#exit	Exit Router mode, and return to Configure mode.

Validation

show ip ospf multi-area-adjacencies, show ip ospf neighbor, show ip ospf route, show ip route

R2

```
R2#show ip ospf multi-area-adjacencies
  Multi-area-adjacency link on interface eth1 to neighbor 2.2.2.2
  Internet Address 2.2.2.1/24, Area 0.0.0.1, MTU 1500
  Process ID 1, Router ID 10.12.26.89, Network Type POINTTOPOINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 16 sent 53, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 10 sent 5
  LS-Ack received 3 sent 9, Discarded 0

R2#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
```

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
10.12.26.88       1   Full/DR         00:00:35    1.1.1.1         eth2
0
10.12.26.90       1   Full/Backup     00:00:33    2.2.2.2         eth1
0
10.12.26.90         1    Full/ -           00:00:35     2.2.2.2          eth1

R2#show ip ospf route

OSPF process 1:
Codes: C - connected, D - Discard, 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

C  1.1.1.0/24 [1] is directly connected, eth2, Area 0.0.0.1
C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  3.3.3.0/24 [2] via 2.2.2.2, eth1, Area 0.0.0.1


R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
C       1.1.1.0/24 is directly connected, eth2
C       2.2.2.0/24 is directly connected, eth1
O       3.3.3.0/24 [110/2] via 2.2.2.2, eth1, 00:05:44
C       10.12.26.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo

Gateway of last resort is not set
```

R3

```
R3#show ip ospf multi-area-adjacencies
  Multi-area-adjacency link on interface eth1 to neighbor 2.2.2.1
  Internet Address 2.2.2.2/24, Area 0.0.0.1, MTU 1500
  Process ID 1, Router ID 10.12.26.90, Network Type POINTTOPOINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 41 sent 41, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 5 sent 10
  LS-Ack received 8 sent 3, Discarded 0

R3#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
```

```
10.12.26.89      1   Full/DR         00:00:39    2.2.2.1       eth1
0
10.12.26.92      1   Full/Backup     00:00:36    3.3.3.2       eth2
0
10.12.26.89       1    Full/ -          00:00:30    2.2.2.1         eth1
R3#sh ip ospf route

OSPF process 1:
Codes: C - connected, D - Discard, 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

O  1.1.1.0/24 [2] via 2.2.2.1, eth1, Area 0.0.0.1
C  2.2.2.0/24 [1] is directly connected, eth1, Area 0.0.0.0
C  3.3.3.0/24 [1] is directly connected, eth2, Area 0.0.0.1

R3#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
O       1.1.1.0/24 [110/2] via 2.2.2.1, eth1, 00:07:31
C       2.2.2.0/24 is directly connected, eth1
C       3.3.3.0/24 is directly connected, eth2
C       10.12.26.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo

Gateway of last resort is not set
```

LSA Throttling

This section contains basic OSPF LSA throttling configuration examples.

The OSPF Link-State Advertisement (LSA) throttling feature provides a mechanism to dynamically slow down link-state advertisement (LSA) updates in OSPF during times of network instability. It also allows faster OSPF convergence by providing LSA rate limiting in milliseconds, when the network is stable.

How OSPF LSA Throttling Works

The `timers throttle lsa all` command controls the generation (sending) of LSAs. The first LSA is always generated immediately upon an OSPF topology change, and the next LSA generated is controlled by the minimum start interval. The subsequent LSAs generated for the same LSA are rate-limited until the maximum interval is reached. The “same LSA” is defined as an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID.

The `timers lsa arrival` command controls the minimum interval for accepting the same LSA. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped. It is recommended that the arrival interval be less than or equal to the hold-time interval of the `timers throttle lsa all` command.

Topology

The diagram shows the minimum configuration required to enable OSPF LSA Throttling Timers feature. R1 and R2 are two routers in Area 0 connecting to network 10.10.10.0/24.

Note: Configure one interface so that it belongs to only one area. It is possible, however, to configure different interfaces on a router to belong to different areas.

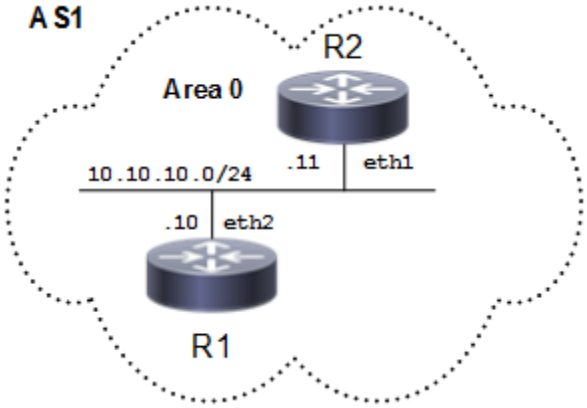

Figure 1-39: Basic OSPF Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the interface loopback to configure.
`(config-if)#ip address 1.1.1.1/32`	Configure the ip address (1.1.1.1) to interface loopback.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (`1`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).
`(config-router)#network 1.1.1.1/32 area 0`	Define the interface (`1.1.1.1/32`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).
`(config-router)#timers throttle lsa all 10000 20000 45000`	Configure LSA Throttling timers (Starting interval: <0-600000>, Min Hold Interval: <1-600000> and Max Wait Interval:< 1-600000>) in milliseconds. The Default value for corresponding timers are: Starting interval: 0, Min Hold Interval: 5 sec and Max Wait Interval: 5 sec.
`(config-router)#logging monitor 7`	Enable logging monitor globally.
`(config-router)#logging level ospf 7`	Enable logging level ospf globally.
`(config-router)#end`	Exit router mode

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID `(1)`. The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface `(10.10.10.0/24)` on which OSPF runs, and associate the area ID `(0)` with the interface.

Validation

R1

Check the output of `show ip ospf` and verify the initial throttle delay, minimum hold time for LSA throttle and maximum wait time for LSA throttle.

```
#show ip ospf 1
 Routing Process "ospf 1" with ID 1.1.1.1
 Process uptime is 11 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 10 secs 0 msecs
 Minimum hold time for LSA throttle 20 secs 0 msecs
 Maximum wait time for LSA throttle 45 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 4
 Number of LSA received 4
 Number of areas attached to this router: 1
MemPool - struct ospf lsa             : (0-8) | Total (8/100000) blk_size:160
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 2(2)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:10:12.807 ago
        SPF algorithm executed 5 times
        Number of LSA 7. Checksum 0x02c480
 Dste Staus: Disabled


#show ip ospf neighbor
```

```
Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID  Pri State       Dead Time  Address     Interface Instance ID
3.1.1.1      1   Full/Backup 00:00:34   10.10.10.11   eth1       0


#debug ospf database-timer rate-limit

#show debugging ospf
OSPF debugging status:
  OSPF rate limit timer events debugging is on
```

Here, we administratively shutdown and then bring up the loopback interface to generate Rate Limit Timer events for OSPF debugging to capture.

```
(config)#int lo
(config-if)#shutdown
2019 Mar 29 16:32:36.838 : OcNOS : OSPF : NOTIF : [OSPF_OPR_LINK_DOWN_4]:
Received Link down for interface: lo
2019 Mar 29 16:32:36.838 : OcNOS : OSPF : INFO : Starting Rate Limit Timer for
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: with 10000 msec delay
2019 Mar 29 16:32:36.838 : OcNOS : OSPF : NOTIF : [OSPF_OPR_STATE_4]:
[lo:1.1.1.1]: Status change Loopback -> Down

(config-if)#no shutdown
2019 Mar 29 16:32:42.705 : OcNOS : OSPF : NOTIF : [OSPF_OPR_LINK_UP_4]:
Received Link up for  interface: lo
2019 Mar 29 16:32:42.705 : OcNOS : OSPF : NOTIF : [OSPF_OPR_STATE_4]:
[lo:1.1.1.1]: Status change Down -> Loopback
2019 Mar 29 16:32:46.853 : OcNOS : OSPF : INFO : Rate Limit Timer for
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: expired
2019 Mar 29 16:32:46.853 : OcNOS : OSPF : INFO : For Next Instance of
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: generation wait 20000 msec


(config-if)#shutdown
2019 Mar 29 16:32:54.353 : OcNOS : OSPF : NOTIF : [OSPF_OPR_LINK_DOWN_4]:
Received Link down for interface: lo
2019 Mar 29 16:32:54.353 : OcNOS : OSPF : INFO : Starting Rate Limit Timer for
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: with 12499 msec delay
2019 Mar 29 16:32:54.353 : OcNOS : OSPF : NOTIF : [OSPF_OPR_STATE_4]:
[lo:1.1.1.1]: Status change Loopback -> Down

(config-if)#no shutdown
2019 Mar 29 16:32:59.252 : OcNOS : OSPF : NOTIF : [OSPF_OPR_LINK_UP_4]:
Received Link up for  interface: lo
2019 Mar 29 16:32:59.252 : OcNOS : OSPF : NOTIF : [OSPF_OPR_STATE_4]:
[lo:1.1.1.1]: Status change Down -> Loopback
2019 Mar 29 16:33:06.870 : OcNOS : OSPF : INFO : Rate Limit Timer for
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: expired
2019 Mar 29 16:33:06.870 : OcNOS : OSPF : INFO : For Next Instance of
LSA[0.0.0.0:Type1:1.1.1.1:(self)]: generation wait 40000 msec
```

R2

Check the output of "`show ip ospf neighbor`" and verify that OSPF adjacency is up.

```
#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
```

```
Neighbor ID  Pri  State     Dead Time   Address    Interface  Instance ID
1.1.1.1      1    Full/DR   00:00:33    10.10.10.10   eth1       0
```

Check the output of `show ip ospf database` and verify that LSA (router LSA in this example) is updated according to the configured LSA throttling timers configured on its neighbor.

```
#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          373     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           71     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          375     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1          372     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          373     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          372     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          373     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          378     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           76     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          380     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router       Age      Seq#     CkSum Opaque ID
1.0.0.1         3.1.1.1          377     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          378     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          377     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          378     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)
```

```
                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          380     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           78     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          382     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1          379     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          380     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          379     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          380     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          381     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           79     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          383     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1          380     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          381     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          380     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          381     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          382     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           80     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)
```

```
Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          384     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1          381     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          382     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          381     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          382     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1          383     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1           81     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1          385     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1          382     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1          383     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1          382     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1          383     0x80000001 0x566c 8
```

Configure OSPF LSA Arrival Timers

The diagram shows the minimum configuration required to enable OSPF Minimum LSA Arrival Timers feature. R1 and R2 are two routers in Area 0 connecting to network 10.10.10.0/24.

Note: Configure one interface so that it belongs to only one area. It is possible, however, to configure different interfaces on a router to belong to different areas.

Topology

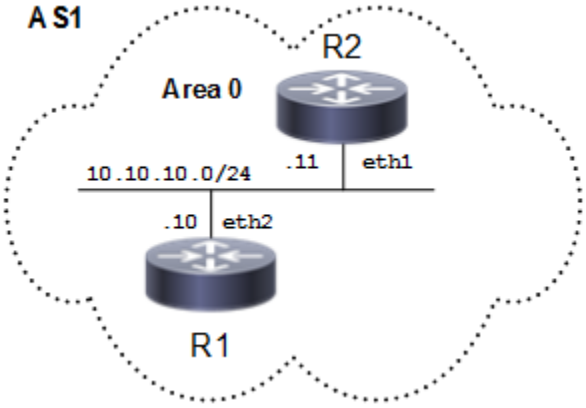

Figure 1-40: Basic OSPF Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the interface loopback to configure.
`(config-if)#ip address 1.1.1.1/32`	Configure the ip address (1.1.1.1) to interface loopback.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (`1`). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 10.10.10.0/24 area 0`	Define the interface (`10.10.10.0/24`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).
`(config-router)#network 1.1.1.1/32 area 0`	Define the interface (`1.1.1.1/32`) on which OSPF runs, and associate the area ID (`0`) with the interface (area ID `0` specifies the backbone area).
`(config-router)#end`	Exit router mode

R2

`#configure terminal`	Enter configure mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID `(1)`. The Process ID should be a unique positive integer identifying the routing process.

`(config-router)#network 10.10.10.0/24 area 0`	Define the interface `(10.10.10.0/24)` on which OSPF runs, and associate the area ID `(0)` with the interface.
`(config-router)#timers lsa arrival 100000`	Configure Minimum LSA Arrival timers (Minimum LSA arrival Interval:< 0-600000>) in milliseconds. The Default value for Minimum LSA Arrival timer is: 1 sec.

Validation

R1

Check the output of `show ip ospf` and verify that the minimum LSA arrival timer by default is set to 1 sec.

```
#show ip ospf
 Routing Process "ospf 1" with ID 1.1.1.1
 Process uptime is 11 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 10 secs 0 msecs
 Minimum hold time for LSA throttle 20 secs 0 msecs
 Maximum wait time for LSA throttle 45 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 4
 Number of LSA received 4
 Number of areas attached to this router: 1
MemPool - struct ospf lsa             : (0-8) | Total (8/100000) blk_size:160
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 2(2)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:10:12.807 ago
        SPF algorithm executed 5 times
        Number of LSA 7. Checksum 0x02c480
 Dste Staus: Disabled

#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID    Pri   State          Dead Time   Address         Interface
Instance ID
3.1.1.1          1   Full/Backup      00:00:34    10.10.10.11     eth1
```

R2

Check the output of `show ip ospf` and verify that the minimum LSA arrival timer is set to 100 sec.

```
#show ip ospf
 Routing Process "ospf 1" with ID 3.1.1.1
 Process uptime is 23 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 100 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 3
 Number of LSA received 10
 Number of areas attached to this router: 1
MemPool - struct ospf lsa             : (0-9) | Total (9/100000) blk_size:160
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:22:12.911 ago
        SPF algorithm executed 4 times
        Number of LSA 7. Checksum 0x02c281
 Dste Staus: Disabled
```

Check the output of show ip ospf neighbor and verify that OSPF adjacency is up.

```
#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
1.1.1.1           1   Full/DR         00:00:35    10.10.10.10     eth1
0
```

Check the output of “show ip ospf database” and verify that LSA is accepted only after a time difference of 100 sec between two consecutive LSAs.

```
#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)
```

```
Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1         1131     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          829     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1         1133     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1         1130     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1131     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1130     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1131     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1         1132     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          831     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1         1134     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1         1131     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1132     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1131     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1132     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1         1133     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          831     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1         1135     0x80000001 0x18e5
```

```
                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#       CkSum  Opaque ID
1.0.0.1         3.1.1.1         1132     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1133     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1132     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1133     0x80000001 0x566c 8


#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum  Link count
3.1.1.1         3.1.1.1         1134     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          832     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum
10.10.10.10     1.1.1.1         1136     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#       CkSum  Opaque ID
1.0.0.1         3.1.1.1         1133     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1134     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1133     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1134     0x80000001 0x566c 8
#
#show ip ospf database

            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum  Link count
3.1.1.1         3.1.1.1         1135     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          834     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum
10.10.10.10     1.1.1.1         1137     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#       CkSum  Opaque ID
1.0.0.1         3.1.1.1         1134     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1135     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1134     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1135     0x80000001 0x566c 8


#show ip ospf database
```

```
            OSPF Router with ID (3.1.1.1) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum  Link count
3.1.1.1         3.1.1.1         1136     0x80000004 0xc60c 1
1.1.1.1         1.1.1.1          834     0x80000008 0xb9f2 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#      CkSum
10.10.10.10     1.1.1.1         1138     0x80000001 0x18e5

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#      CkSum  Opaque ID
1.0.0.1         3.1.1.1         1135     0x80000001 0x2cf6 1
1.0.0.1         1.1.1.1         1136     0x80000001 0x2af6 1
1.0.0.8         3.1.1.1         1135     0x80000001 0x7d45 8
1.0.0.8         1.1.1.1         1136     0x80000001 0x566c 8
```

Loop-Free Alternate Fast Reroute

This section contains basic OSPF Loop-Free Alternate Fast Reroute (LFA-FRR) configuration examples.

Overview

The goal of (LFA-FRR) is to reduce failure reaction time to 10s of milliseconds by using a pre-computed alternate next-hop in the event that the currently selected primary next-hop fails, so that the alternate can be rapidly used when the failure is detected. A network with this feature experiences less traffic loss and less micro-looping of packets than a network without LFA-FRR.

After enabling LFA-FRR on routers, routers calculate a backup path for each primary path to reach the destination.The backup path is calculated based on the attributes such as node protecting, link protecting, broadcast-link protecting and secondary path.

Note: RIB errors occur when the hardware profile cannot handle LFA-ECMP. The error prompts users to alter or reverse the configurations.

Topology

The diagram shows the configuration required to enable the OSPF LFA feature.

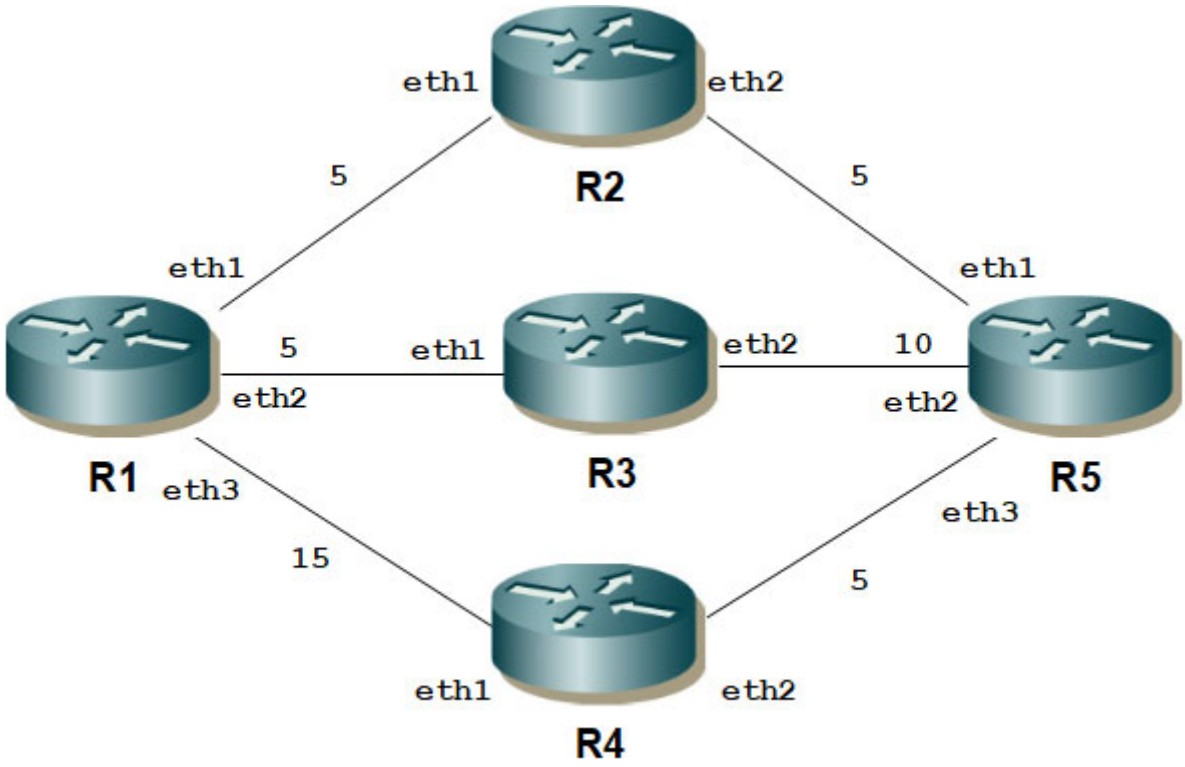

Figure 1-41: Figure 4-13: Basic OSPF-LFA Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#int eth1`	Enter interface mode.
`(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth2`	Enter interface mode.
`(config-if)#ip address 20.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#int eth3`	Enter interface mode.
`(config-if)#ip address 30.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the routing process and specify the Process ID (1).
`(config-router)#network 10.1.1.0/24 area 0`	Define the interface (10.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 20.1.1.0/24 area 0`	Define the interface (20.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

© 2025 IP Infusion Inc. Proprietary

(config-router)#network 30.1.1.0/24 area 0	Define the interface (30.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#fast-reroute keep-all- paths	Configure LFA-FRR to calculate the available backup path.
(config-router)#end	Exit router mode.

R2

#configure terminal	Enter configure mode.
(config)#int eth1	Enter interface mode.
(config-if)#ip address 10.1.1.2/24	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#int eth2	Enter interface mode.
(config-if)#ip address 40.1.1.1/24	Configure the IP address of the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
(config-router)#network 10.1.1.0/24 area 0	Define the interface (20.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#network 40.1.1.0/24 area 0	Define the interface (40.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#end	Exit router mode.

R3

#configure terminal	Enter configure mode.
(config)#int eth1	Enter interface mode.
(config-if)#ip address 20.1.1.2/24	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#int eth2	Enter interface mode.
(config-if)#ip address 50.1.1.1/24	Configure the IP address of the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
(config-router)#network 20.1.1.0/24 area 0	Define the interface (20.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#network 50.1.1.0/24 area 0	Define the interface (50.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#end	Exit router mode.

R4

`#configure terminal`	Enter configure mode.
`(config)#int eth1`	Enter interface mode.
`(config-if)#ip address 30.1.1.2/24`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth2`	Enter interface mode.
`(config-if)#ip address 60.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 30.1.1.0/24 area 0`	Define the interface (30.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 60.1.1.0/24 area 0`	Define the interface (60.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#end`	Exit router mode.

© 2025 IP Infusion Inc. Proprietary

R5

`#configure terminal`	Enter configure mode.
`(config)#int eth1`	Enter interface mode.
`(config-if)#ip address 40.1.1.2/24`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode.
`(config)#int eth2`	Enter interface mode.
`(config-if)#ip address 50.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#int eth3`	Enter interface mode.
`(config-if)#ip address 60.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#network 40.1.1.0/24 area 0`	Define the interface (30.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 50.1.1.0/24 area 0`	Define the interface (30.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 60.1.1.0/24 area 0`	Define the interface (30.1.1.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#end`	Exit router mode.

Validation

R1

Check OSPF neighborship.

```
#show ip ospf neighbor
OSPF Process 100 VRF (default)
Neighbor ID     Pri   State           Dead Time   Address  Interface Intance ID
2.2.2.2           1   Full/DR         00:00:33    10.1.1.2     eth1  0
3.3.3.3           1   Full/DR         00:00:33    20.1.1.2     eth2  0
4.4.4.4           1   Full/DR         00:00:39    30.1.1.2     eth2  0
#
```

Check the OSPF route installation and LFA-FRR backup path for the primary path.

```
#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
```

```
area
* - candidate default

IP Route Table for VRF "default"
C     10.1.1.0/24 is directly connected, eth1
C     20.1.1.0/24 is directly connected, eth2
C 30.1.1.0/24 is directly connected, eth3
O     40.1.1.0/24 [110/10] via 10.1.1.2, eth1, 00:16:43
O     50.1.1.0/24 [110/15] via 20.1.1.2, eth2, 00:16:43
O IA  60.1.1.0/24 [110/15] via 10.1.1.2, eth1, 00:16:43
O E2  70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 00:16:43
O E2 80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 00:16:43
C 127.0.0.0/8 is directly connected, lo
C 192.168.100.0/24 is directly connected, eth0
Gateway of last resort is not set
```

Not mandatory that for all primary path, there exists an LFA backup path only if inequality equation satisfies according to attributes configured on routers, backup path will be calculated.

```
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
       * - candidate default

IP Route Table for VRF "default"
O       50.1.1.0/24 [110/15] via 20.1.1.2, eth2, 00:00:34
                    [FRR-NH] via 10.1.1.2, eth1

O       60.1.1.0/24 [110/15] via 10.1.1.2, eth1, 00:00:34
                    [FRR-NH] via 20.1.1.2, eth2

O       70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 00:02:27
                    [FRR-NH] via 10.1.1.2, eth1

O       80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 20.1.1.2, eth2
```

Not mandatory that for all primary path, there exists an LFA backup path only if inequality equation satisfies according to attributes configured on routers, backup path will be calculated.

To prohibit an interface from being used as a repair path, disable fast reroute calculation on the interface.

```
(config)#int eth3
(config-if)#ip ospf fast-reroute per-prefix candidate disable
(config-if)#end
```

Verify that the eth3 interface is not used for backup path calculation.

```
#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

IP Route Table for VRF "default"
O       10.1.1.0/24 [110/10] via 20.1.1.1, eth1, 00:34:04
C       20.1.1.0/24 is directly connected, eth1
O       30.1.1.0/24 [110/20] via 20.1.1.1, eth1, 00:34:04
O       40.1.1.0/24 [110/15] via 20.1.1.1, eth1, 00:34:04
                    [110/15] via 50.1.1.2, eth2, 00:34:04
C       50.1.1.0/24 is directly connected, eth2
O IA    60.1.1.0/24 [110/15] via 50.1.1.2, eth2, 01:08:29
O E2    70.1.1.0/24 [110/20] via 20.1.1.1, eth1, 00:34:03
O E2    80.1.1.0/24 [110/20] via 50.1.1.2, eth2, 01:11:17
C       127.0.0.0/8 is directly connected, lo
C       192.168.100.0/24 is directly connected, eth0
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
       * - candidate default

IP Route Table for VRF "default"
O       10.1.1.0/24 [110/10] via 20.1.1.1, eth1, 00:00:34
                    [FRR-NH] via 50.1.1.2, eth2

O       30.1.1.0/24 [110/20] via 20.1.1.1, eth1, 00:00:34
                    [FRR-NH] via 50.1.1.2, eth2

O       60.1.1.0/24 [110/15] via 50.1.1.2, eth2, 00:02:27
                    [FRR-NH] via 20.1.1.1, eth1

O       70.1.1.0/24 [110/20] via 20.1.1.1, eth1, 00:02:27
                    [FRR-NH] via 50.1.1.2, eth2

O       80.1.1.0/24 [110/20] via 50.1.1.2, eth2, 00:02:27
```

Note: Now the LFA backup paths have been changed, eth3 is not used.

LFA Tie-Breaker

Based on the index values configured, if inequalities are satisfied, protections will be provided:

- Lower the index will have the highest priority, the path which provides protection with highest priority will be selected. If there are multiple paths providing the highest priority protection, then we will check which path provides the protection which has 2nd highest priority and so on.
- If all the paths provide same priority, then the LFA route is chosen on the basis of path cost.
- If none of the paths provides the protection with highest priority, then we will see which path provides the 2nd highest priority and so on.

```
config)#router ospf 100
(config-router)#fast-reroute tie-break ?
broadcast-interface-disjoint  Prefer broadcast link protecting backup path

                              *Default value is 70

downstream-path               Prefer backup path from downstream

                              *Default value is 90

interface-disjoint            Prefer link protecting backup path

                              *Default value is 60

node-protecting               Prefer node protecting backup path

                              *Default value is 30

primary-path                  Prefer backup path from ECMP set

                              *Default value is 20

secondary-path                Prefer non-ECMP backup path

                              *Default value is 255
(config-router)#fast-reroute tie-break broadcast-interface-disjoint index 1
(config-router)#fast-reroute tie-break node-protecting index 2
```

Verify show ip route and show ip route fast-reroute for backup path calculated according to attributes configured above.

```
#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP 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
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
* - candidate default

IP Route Table for VRF "default"
C     10.1.1.0/24 is directly connected, eth1
C     20.1.1.0/24 is directly connected, eth2
```

```
C 30.1.1.0/24 is directly connected, eth3
O     40.1.1.0/24 [110/10] via 10.1.1.2, eth1, 01:07:26
O     50.1.1.0/24 [110/15] via 20.1.1.2, eth2, 01:07:26
O IA  60.1.1.0/24 [110/15] via 10.1.1.2, eth1, 01:07:26
O E2  70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 01:07:26
O E2 80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 01:07:26
C 127.0.0.0/8 is directly connected, lo
C 192.168.100.0/24 is directly connected, eth0

Gateway of last resort is not set
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
       * - candidate default

IP Route Table for VRF "default"

O       50.1.1.0/24 [110/15] via 20.1.1.2, eth2, 00:00:34
                    [FRR-NH] via 10.1.1.2, eth1

O       60.1.1.0/24 [110/15] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 20.1.1.2, eth2

O       70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 00:02:27
                    [FRR-NH] via 10.1.1.2, eth1

O       80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 20.1.1.2, eth2

#show ip ospf route fast-reroute

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, ID - Interface Disjoint,
       NP - Node Protecting, BID - Broadcast Interface Disjoint

O  50.1.1.0/24 [15] via 20.1.1.1, eth2, Area 0.0.0.0
                       Backup path:
                        via 10.1.1.2, eth1, Area 0.0.0.0
                       Attributes: Metric: [20] ,LP ,NP, BP
O  60.1.1.0/24 [15] via 10.1.1.2, eth1, Area 0.0.0.0
                       Backup path:
                        via 20.1.1.2, eth2, Area 0.0.0.0
                       Attributes: Metric: [15] ,LP ,NP,BP
```

```
O  70.1.1.0/24 [20] via 30.1.1.2, eth3, Area 0.0.0.0
                       Backup path:
                        via 10.1.1.2, eth1, Area 0.0.0.0
                       Attributes: Metric: [20] ,LP ,NP,BP
O  80.1.1.0/24 [20] via 10.1.1.2, eth1, Area 0.0.0.0
                       Backup path:
                        via 20.1.1.2, eth2, Area 0.0.0.0
                       Attributes: Metric: [20] ,LP,NP,BP
```

LFA Termination

A router MUST limit the amount of time an alternate next-hop is used after the primary next-hop has become unavailable. This ensures that the router will start using the new primary next-hops.

LFA termination avoids a micro looping in topology, when particular network goes down, LFA backup path will be installed and if termination interval is configured, LFA backup will be still used till the interval and it is used in order to verify new primary path is loop free.

R1

Configure termination interval on R1 in router mode:

`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config-router)#fast-reroute terminate-hold-on interval 100000`	Configure LFA termination interval
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

If you check "show ip ospf" you can see the configured termination-hold on interval value along with ospf output:

```
#show ip ospf
IPFRR per-prefix tiebreakers:
    Name                                Index
    Primary path                        20
    Node Protecting                     30
    Interface disjoint                  60
    Broadcast interface disjoint        70
    Secondary path                      255
 LFA termination hold-on timer : 100 secs 0 msecs

#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
       * - candidate default

IP Route Table for VRF "default"
```

```
O       50.1.1.0/24 [110/15] via 20.1.1.2, eth2, 00:00:34
                    [FRR-NH] via 10.1.1.2, eth1

O       60.1.1.0/24 [110/15 via 10.1.1.2, eth1, 00:00:34
                    [FRR-NH] via 20.1.1.2, eth2

O       70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 00:02:27
                    [FRR-NH] via 10.1.1.2, eth1

O       80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 20.1.1.2, eth2
```

Shut down one of the primary nexthops, here eth2 of rtr1:

`(config)#interface eth2`	Enter interface mode.
`(config-if)#shutdown`	Shutdown the interface
`(config-if)#exit`	Exit interface mode
`(config-if)#exit`	Exit interface mode.

```
#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p -
stale info
       * - candidate default

IP Route Table for VRF "default"
O       50.1.1.0/24 [110/15] via 20.1.1.2, eth1, 00:00:34  <<<eth1 which was back-up
path before got  installed as new primary path
                    [FRR-NH] via 30.1.1.2, eth3

O       60.1.1.0/24 [110/15] via 10.1.1.2, eth1, 00:00:34
                    [FRR-NH] via 30.1.1.2, eth3

O       70.1.1.0/24 [110/20] via 30.1.1.2, eth3, 00:02:27
                    [FRR-NH] via 10.1.1.2, eth1

O       80.1.1.0/24 [110/20] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 30.1.1.2, eth3
```

SNMP Support for Multiple Instance of OSPF Configuration

Overview

SNMP operation by default are tied to a specific OID which is unique. However protocol like OSPF can have multiple instances, and have different values of same parameters for different OSPF instances. To be able to support SNMP for each of these instances, it is needed that each instance of the protocol has its own instance of the MIBs. It is aimed to achieve that with mapping each instance to a context. Each context will point to a different copy of the same OID for the protocol.

Topology

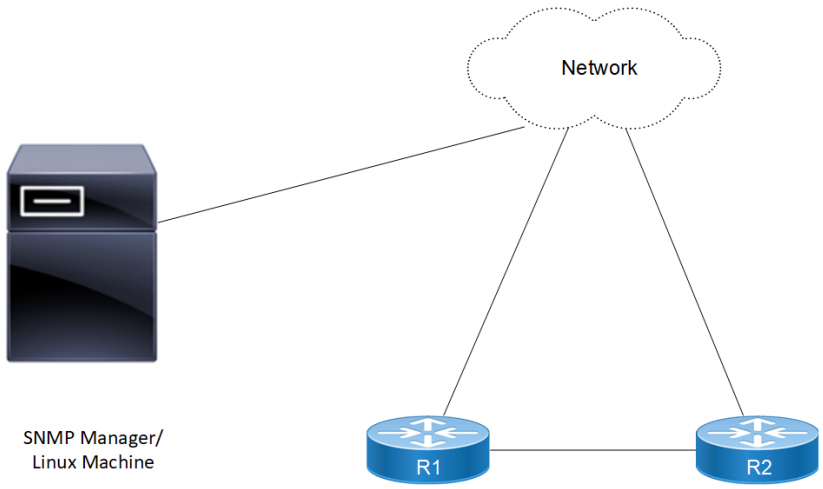

Figure 1-42: SNMP OSPF instance

Enable SNMP and create SNMP Context & Group for OSPF Instances on a Router with SNMPv2

In this example, routers R1 & R2 are in Area 0, and all run OSPF. SNMPv2 user is created and Mapping of user with group and context for SNMPwalk /SNMP get operation on context.

R1

`#configure terminal`	Enter configure mode.
`(config)#snmp-server enable snmp vrf management`	Use this command to start the SNMP agent.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.
`#configure terminal`	Enter configure mode.
`(config)#snmp-server context context1 vrf management`	Creates SNMP Context with Context name.
`(config)#snmp-server group group1 version 2c context context1 vrf management`	Creates SNMP group with Group name and for specific context in SNMP v2 version.

(config)#snmp-server user user1 group1 vrf management	Creates SNMP User in SNMPv2 and attach user into a group
(config)#snmp-server community cm1 vrf management	Set community string as "cm1"
(config)#snmp-server community-map cm1 context context1 user user1 vrf management	Creates Community map SNMPv2 with community name mapping user with a context.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 94.94.94.3/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 100	Configure an OSPF instance with an instance ID of 100.
(config-router)#router-id 21.21.21.21	Configure the router ID to use on this instance.
(config-router)#network 94.94.94.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context1	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

#configure terminal	Enter configure mode.
(config)#snmp-server enable snmp vrf management	Use this command to start the SNMP agent.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
#configure terminal	Enter configure mode.
(config)#snmp-server context context1 vrf management	Creates SNMP Context with Context name.
(config)#snmp-server group group1 version 2c context context1 vrf management	Creates SNMP group with Group name and for specific context in SNMP v2 version.
(config)#snmp-server user user1 group1 vrf management	Creates SNMP User in SNMPv2 and attach user into a group
(config)#snmp-server community cm1 vrf management	Set community string as "cm1"
(config)#snmp-server community-map cm1 context context1 user user1 vrf management	Creates Community map SNMPv2 with community name mapping user with a context.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 94.94.94.2/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.

(config)#router ospf 100	Configure an OSPF instance with an instance ID of 100.
(config-router)#router-id 23.23.23.23	Configure the router ID to use on this instance.
(config-router)#network 94.94.94.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context1	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#sh running-config
!
! Software version: EC_AS7326-56X-OcNOS-5.1.194-DC-IPBASE-S0-P0 04/20/2022 18:27:17
!
!Last configuration change at 12:56:57 UTC Fri Apr 22 2022 by ocnos
!
no service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
ip vrf management
!
hostname R1
!
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server group group1 version 2c context context1 vrf management
snmp-server user user1 group1 vrf management
snmp-server community cm1 vrf management
snmp-server community-map cm1 context context1 user user1 vrf management

!
interface ce1
!
interface ce2
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
```

```
 ip address 127.0.0.1/8
 ip address 21.21.21.21/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
ip address 94.94.94.3/24
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
router ospf 100
 ospf router-id 21.21.21.21
 snmp context-name context1
 network 21.21.21.21/32 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
!
end

R1#

R1#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
23.23.23.23       1   Full/ -         00:00:37    94.94.94.2      xe1

R1#sh snmp context

-----------------------------------------------------------------
context                                          groups
```

```
------------------------------------------------------------------------
context1                                        group1
R1#
```

R2

```
R2#sh running-config
!
! Software version: EC_AS5912-54X-OcNOS-5.1.194-SP-IPBASE-S0-P0 04/20/2022 18:28:57
!
!Last configuration change at 13:01:05 UTC Fri Apr 22 2022 by root
!
no service password-encryption
!
logging console 5
logging level all 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hostname R2
bridge 1 protocol ieee vlan-bridge
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server group group1 version 2c context context1 vrf management
snmp-server user user1 group1 vrf management
snmp-server community cm1 vrf management
snmp-server community-map cm1 context context1 user user1 vrf management
feature ntp vrf management
ntp enable vrf management
!
interface ce49
!
interface ce50
!
interface ce51
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 23.23.23.23/32 secondary
 ipv6 address ::1/128
!
interface lo.management
```

```
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
ip address 94.94.94.2/24
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
router ospf 100
 ospf router-id 23.23.23.23
 snmp context-name context1
 network 23.23.23.23/32 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
!
end

R2#
R2#sh snmp context

-----------------------------------------------------------------------
context                                          groups
-----------------------------------------------------------------------
context1                                         group1
R2#
```

SNMP WALK Command

Perform snmpwalk as mentioned below with IPv4 address using SNMPv2

```
snmpwalk -v2c -c cm1 10.12.86.116 .1.3


[root@localhost ~]# snmpwalk -v2c -c cm1 10.12.86.116 .1.3
OSPF-MIB::ospfRouterId.0 = IpAddress: 23.23.23.23
OSPF-MIB::ospfAdminStat.0 = INTEGER: enabled(1)
OSPF-MIB::ospfVersionNumber.0 = INTEGER: version2(2)
```

```
OSPF-MIB::ospfExternLsaCount.0 = Gauge32: 0
OSPF-MIB::ospfExternLsaCksumSum.0 = INTEGER: 0
OSPF-MIB::ospfTOSSupport.0 = INTEGER: false(2)
OSPF-MIB::ospfOriginateNewLsas.0 = Counter32: 3
OSPF-MIB::ospfRxNewLsas.0 = Counter32: 14
OSPF-MIB::ospfExtLsdbLimit.0 = INTEGER: -1
OSPF-MIB::ospfMulticastExtensions.0 = INTEGER: 0
OSPF-MIB::ospfExitOverflowInterval.0 = INTEGER: 0
OSPF-MIB::ospfDemandExtensions.0 = INTEGER: false(2)
OSPF-MIB::ospfRFC1583Compatibility.0 = INTEGER: false(2)
OSPF-MIB::ospfOpaqueLsaSupport.0 = INTEGER: true(1)
OSPF-MIB::ospfReferenceBandwidth.0 = Gauge32: 100000 kilobits per second
OSPF-MIB::ospfRestartSupport.0 = INTEGER: plannedAndUnplanned(3)
OSPF-MIB::ospfRestartInterval.0 = INTEGER: 120 seconds
OSPF-MIB::ospfRestartStrictLsaChecking.0 = INTEGER: true(1)
OSPF-MIB::ospfRestartStatus.0 = INTEGER: notRestarting(1)
OSPF-MIB::ospfRestartAge.0 = Gauge32: 0 seconds
OSPF-MIB::ospfRestartExitReason.0 = INTEGER: none(1)
OSPF-MIB::ospfAsLsaCount.0 = Gauge32: 0
OSPF-MIB::ospfAsLsaCksumSum.0 = Gauge32: 0
OSPF-MIB::ospfStubRouterSupport.0 = INTEGER: false(2)
OSPF-MIB::ospfStubRouterAdvertisement.0 = INTEGER: doNotAdvertise(1)
OSPF-MIB::ospfDiscontinuityTime.0 = Timeticks: (0) 0:00:00.00
OSPF-MIB::ospfAreaId.0.0.0.0 = IpAddress: 0.0.0.0
OSPF-MIB::ospfAuthType.0.0.0.0 = INTEGER: none(0)
OSPF-MIB::ospfImportAsExtern.0.0.0.0 = INTEGER: importExternal(1)
OSPF-MIB::ospfSpfRuns.0.0.0.0 = Counter32: 5
OSPF-MIB::ospfAreaBdrRtrCount.0.0.0.0 = Gauge32: 0
OSPF-MIB::ospfAsBdrRtrCount.0.0.0.0 = Gauge32: 0
OSPF-MIB::ospfAreaLsaCount.0.0.0.0 = Gauge32: 6
OSPF-MIB::ospfAreaLsaCksumSum.0.0.0.0 = INTEGER: 199510
OSPF-MIB::ospfAreaSummary.0.0.0.0 = INTEGER: sendAreaSummary(2)
OSPF-MIB::ospfAreaStatus.0.0.0.0 = INTEGER: active(1)
OSPF-MIB::ospfAreaNssaTranslatorRole.0.0.0.0 = INTEGER: 0
OSPF-MIB::ospfAreaNssaTranslatorState.0.0.0.0 = INTEGER: 0
OSPF-MIB::ospfAreaNssaTranslatorStabilityInterval.0.0.0.0 = INTEGER: 40 seconds
OSPF-MIB::ospfAreaLsaCountNumber.0.0.0.0 = Gauge32: 6
OSPF-MIB::ospfAreaLsaCountNumber.0.0.0.0 = No more variables left in this MIB View (It
is past the end of the MIB tree)
[root@localhost ~]#
```

Perform snmpwalk as mentioned below with IPv4 address using SNMPv2 for R2

```
snmpwalk -v2c -c cm1 10.12.86.111 .1.3
```

Enable SNMP and create SNMP Context & Group for OSPF Instances on a Router with SNMPv3 Configuration

In this example, routers R1 & R2 are in Area 0, and all run OSPF. SNMPv3 user is created and Mapping of user with group and context for SNMPwalk /SNMP get operation on context.

R1

`#configure terminal`	Enter configure mode.
`(config)#snmp-server enable snmp vrf management`	Use this command to start the SNMP agent.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.
`#configure terminal`	Enter configure mode.
`(config)#snmp-server context context1 vrf management`	Creates SNMP Context with Context name.
`(config)#snmp-server group group2 version 3 auth context context1 vrf management`	Creates SNMP group with Group name and for specific context in SNMP v3 version.
`(config)#snmp-server user user2 group2 auth md5 password vrf management`	Creates SNMP User in SNMPv3 and attach user into a group
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.
`(config)#interface xe1`	Enter interface mode for xe1.
`(config-if)#ip address 94.94.94.3/24`	Specify the IP address of the interface.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 100`	Configure an OSPF instance with an instance ID of 100.
`(config-router)#router-id 21.21.21.21`	Configure the router ID to use on this instance.
`(config-router)#network 94.94.94.0/24 area 0`	Advertise the network with the area ID.
`(config-router)#snmp context-name context1`	SNMP Context is mapped with OSPF Instance
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#snmp-server enable snmp vrf management`	Use this command to start the SNMP agent.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.
`#configure terminal`	Enter configure mode.
`(config)#snmp-server context context1 vrf management`	Creates SNMP Context with Context name.
`(config)#snmp-server group group2 version 3 auth context context1 vrf management`	Creates SNMP group with Group name and for specific context in SNMP v3 version.
`(config)#snmp-server user user2 group2 auth md5 password vrf management`	Creates SNMP User in SNMPv2 and attach user into a group
`(config)#commit`	Commit the candidate configuration to the running configuration

(config)#exit	Exit configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 94.94.94.2/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 100	Configure an OSPF instance with an instance ID of 100.
(config-router)#router-id 23.23.23.23	Configure the router ID to use on this instance.
(config-router)#network 94.94.94.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context1	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#sh running-config
!
! Software version: EC_AS7326-56X-OcNOS-5.1.194-DC-IPBASE-S0-P0 04/20/2022 18:27:1
7
!
!Last configuration change at 14:06:07 UTC Fri Apr 22 2022 by root
!
no service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
ip vrf management
!
hostname R1
ip name-server vrf management 10.12.3.23
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server group grp2 version 3 auth context context2 vrf management
snmp-server user user2 grp2 auth md5 encrypt 0x2eaaa9043312c907 vrf management
feature ntp vrf management
ntp enable vrf management
feature rsyslog vrf management
!
interface ce49
!
```

```
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
!
interface ce54
!
interface ce55
!
interface ce56
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 21.21.21.21/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
 ip address 94.94.94.3/24
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
router ospf 100
```

```
 ospf router-id 21.21.21.21
 snmp context-name context1
 network 21.21.21.21/32 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
!
line console 0
 exec-timeout 0 0
line vty 0 871
 exec-timeout 0 0
 privilege level 16
!
!
end

R1#
```

R2

```
R2#  sh run
!
! Software version: EC_AS5912-54X-OcNOS-5.1.194-SP-IPBASE-S0-P0 04/20/2022 18:28:5
7
!
!Last configuration change at 14:28:48 UTC Fri Apr 22 2022 by root
!
no service password-encryption
!
logging console 5
logging level all 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hostname R2
no ip domain-lookup
ip domain-lookup vrf management
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server group group2 version 3 auth context context1 vrf management
snmp-server user user2 group2 auth md5 encrypt 0x2eaaa9043312c907 vrf management
feature ntp vrf management
ntp enable vrf management
feature rsyslog vrf management
!
interface ce49
!
interface ce50
```

```
!
interface ce51
!
interface ce52
!
interface ce53
!
interface ce54
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 23.23.23.23/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
 ip address 94.94.94.2/24
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
router ospf 100
 ospf router-id 23.23.23.23
 bfd all-interfaces
 network 23.23.23.23/32 area 0.0.0.0
 network 91.91.91.0/24 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
```

```
 network 96.96.96.0/24 area 0.0.0.0
!
!
end

R2#
```

SNMP WALK Command

Perform snmpwalk as mentioned below with IPv4 address using SNMPv3 of R2

```
snmpwalk -v 3 -u user2 -l auth -r 0 -t 10 -n "context1" -a MD5 -A password  10.12.86.116
iso.3.6.1.2.1

[root@localhost ~]# snmpwalk -v 3 -u user2 -l auth -r 0 -t 10 -n "context1" -a MD5 -A
password  10.12.86.116 iso.3.6.1.2.1
OSPF-MIB::ospfRouterId.0 = IpAddress: 23.23.23.23
OSPF-MIB::ospfAdminStat.0 = INTEGER: enabled(1)
OSPF-MIB::ospfVersionNumber.0 = INTEGER: version2(2)
OSPF-MIB::ospfExternLsaCount.0 = Gauge32: 0
OSPF-MIB::ospfExternLsaCksumSum.0 = INTEGER: 0
OSPF-MIB::ospfTOSSupport.0 = INTEGER: false(2)
OSPF-MIB::ospfOriginateNewLsas.0 = Counter32: 3
OSPF-MIB::ospfRxNewLsas.0 = Counter32: 7
OSPF-MIB::ospfExtLsdbLimit.0 = INTEGER: -1
OSPF-MIB::ospfMulticastExtensions.0 = INTEGER: 0
OSPF-MIB::ospfExitOverflowInterval.0 = INTEGER: 0
OSPF-MIB::ospfDemandExtensions.0 = INTEGER: false(2)
OSPF-MIB::ospfRFC1583Compatibility.0 = INTEGER: false(2)
OSPF-MIB::ospfOpaqueLsaSupport.0 = INTEGER: true(1)
OSPF-MIB::ospfReferenceBandwidth.0 = Gauge32: 100000 kilobits per second
OSPF-MIB::ospfRestartSupport.0 = INTEGER: plannedAndUnplanned(3)
OSPF-MIB::ospfRestartInterval.0 = INTEGER: 120 seconds
OSPF-MIB::ospfRestartStrictLsaChecking.0 = INTEGER: true(1)
OSPF-MIB::ospfRestartStatus.0 = INTEGER: notRestarting(1)
OSPF-MIB::ospfRestartAge.0 = Gauge32: 0 seconds
OSPF-MIB::ospfRestartExitReason.0 = INTEGER: none(1)
OSPF-MIB::ospfAsLsaCount.0 = Gauge32: 0
OSPF-MIB::ospfAsLsaCksumSum.0 = Gauge32: 0
OSPF-MIB::ospfStubRouterSupport.0 = INTEGER: false(2)
OSPF-MIB::ospfStubRouterAdvertisement.0 = INTEGER: doNotAdvertise(1)
OSPF-MIB::ospfDiscontinuityTime.0 = Timeticks: (0) 0:00:00.00
OSPF-MIB::ospfAreaId.0.0.0.0 = IpAddress: 0.0.0.0
OSPF-MIB::ospfAuthType.0.0.0.0 = INTEGER: none(0)
OSPF-MIB::ospfAreaLsaCountNumber.0.0.0.0 = No more variables left in this MIB View (It
is past the end of the MIB tree)
[root@localhost ~]#
Perform snmpwalk as mentioned below with IPv4 address using SNMPv3 for R1.
snmpwalk -v 3 -u user2 -l auth -r 0 -t 10 -n "context1" -a MD5 -A password  10.12.86.111
iso.3.6.1.2.1
```

Enable SNMP and create SNMP Context & Group for Multiple OSPF Instances on a Router with SNMPv2/v3

In this example, routers R1, R2 & R3 are in Area 0, and all run OSPF. SNMPv2/v3 user is created and Mapping of user with group and context for SNMPwalk /SNMP get operation on context.

Topology

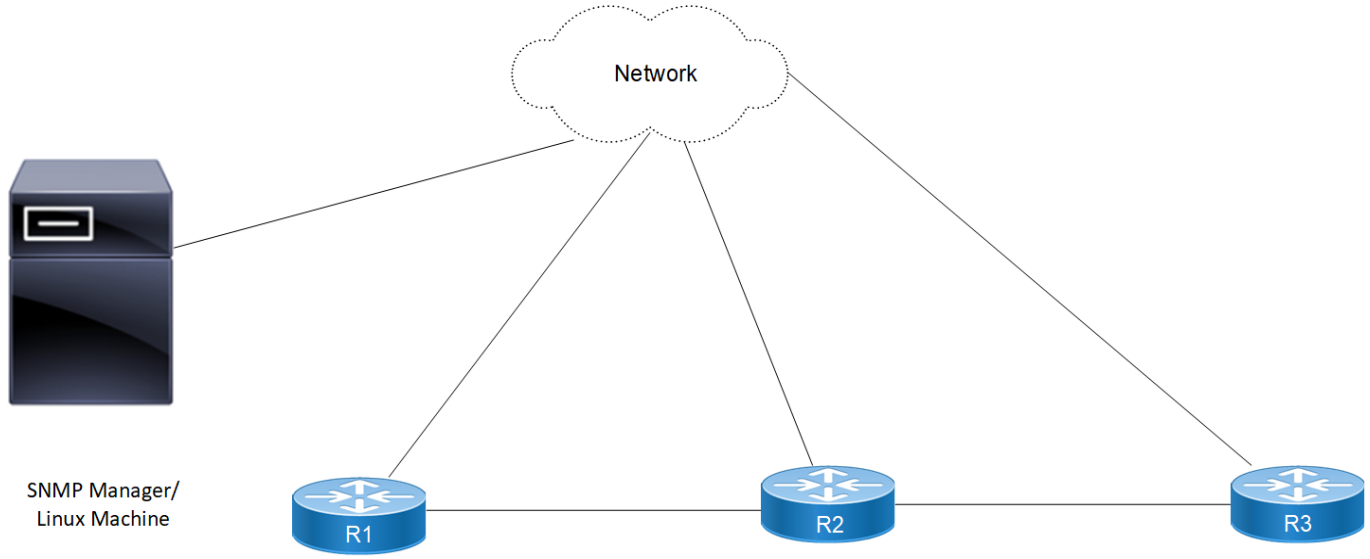

Figure 1-43: SNMP OSPF multiple instance

Configurations

R1

`#configure terminal`	Enter configure mode.
`(config)#snmp-server enable snmp vrf management`	Use this command to start the SNMP agent.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.
`#configure terminal`	Enter configure mode.
`(config)#snmp-server context context1 vrf management`	Creates SNMP Context with Context name.
`(config)#snmp-server group group1 version 2c context context1 vrf management`	Creates SNMP group with Group name and for specific context in SNMP v2 version.
`(config)#snmp-server group group2 version 3 auth context context1 vrf ma`	Creates SNMP group with Group name and for specific context in SNMP v3 version
`(config)#snmp-server user user1 group1 vrf management`	Creates SNMP User in SNMPv2 and attach user into a group
`(config)#snmp-server user user2 group2 auth md5 password vrf  management`	Creates SNMP User in SNMPv3 and attach user into a group

© 2025 IP Infusion Inc. Proprietary

(config)#snmp-server community cm1 vrf management	Set community string as "cm1"
(config)#snmp-server community-map cm1 context context1 user user1 vrf management	Creates Community map SNMPv2 with community name mapping user with a context.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 94.94.94.3/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 100	Configure an OSPF instance with an instance ID of 100.
(config-router)#router-id 21.21.21.21	Configure the router ID to use on this instance.
(config-router)#network 94.94.94.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context1	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

#configure terminal	Enter configure mode.
(config)#snmp-server enable snmp vrf management	Use this command to start the SNMP agent.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
#configure terminal	Enter configure mode.
(config)#snmp-server context context1 vrf management	Creates SNMP Context with Context name "context1".
(config)#snmp-server context context2 vrf management	Creates SNMP Context with Context name"context2".
(config)#snmp-server group group1 version 3 auth context context1 vrf management	Creates SNMP group with Group name and for context 1 in SNMP v3 version.
(config)#snmp-server group group1 version 3 auth context context2 vrf management	Creates SNMP group with Group name and for context 2 in SNMP v3 version.
(config)#snmp-server user user1 group1 auth md5 password vrf management	Creates SNMP User in SNMPv3 and attach user into a group
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
(config)#interface xe1	Enter interface mode for xe1.
(config-if)#ip address 94.94.94.2/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ip address 10.1.2.2/24	Specify the IP address of the interface.

(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 100	Configure an OSPF instance with an instance ID of 100.
(config-router)#router-id 23.23.23.23	Configure the router ID to use on this instance.
(config-router)#network 94.94.94.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context1	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 200	Configure an OSPF instance with an instance ID of 200.
(config-router)#router-id 24.24.24.24	Configure the router ID to use on this instance.
(config-router)#network 10.1.2.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context2	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

R3

#configure terminal	Enter configure mode.
(config)#snmp-server enable snmp vrf management	Use this command to start the SNMP agent.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
#configure terminal	Enter configure mode.
(config)#snmp-server context context2 vrf management	Creates SNMP Context with Context name.
(config)#snmp-server group group1 version 2c context context2 vrf management	Creates SNMP group with Group name and for specific context in SNMP v2 version.
(config)#snmp-server group group2 version 3 auth context context2 vrf management	Creates SNMP group with Group name and for specific context in SNMP v3 version
(config)#snmp-server user user1 group1 vrf management	Creates SNMP User in SNMPv2 and attach user into a group
(config)#snmp-server user user2 group2 auth md5 password vrf management	Creates SNMP User in SNMPv3 and attach user into a group
(config)#snmp-server community cm1 vrf management	Set community string as "cm1"
(config)#snmp-server community-map cm1 context context1 user user1 vrf management	Creates Community map SNMPv2 with community name mapping user with a context.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode.
(config)#interface xe2	Enter interface mode for xe2.
(config-if)#ip address 10.1.2.3/24	Specify the IP address of the interface.
(config-if)#no shutdown	Activate the interface.
(config-if)#exit	Exit interface mode.
(config)#router ospf 200	Configure an OSPF instance with an instance ID of 200.

(config-router)#router-id 24.24.24.24	Configure the router ID to use on this instance.
(config-router)#network 10.1.2.0/24 area 0	Advertise the network with the area ID.
(config-router)#snmp context-name context2	SNMP Context is mapped with OSPF Instance
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show run
!
! Software version: EC_AS7326-56X-OcNOS-5.1.194-DC-IPBASE-S0-P0 04/20/2022 18:27:1
7
!
!Last configuration change at 15:10:23 UTC Fri Apr 22 2022 by root
!
no service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
ip vrf management
!
hostname R1
no ip domain-lookup
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
tfo Disable
errdisable cause stp-bpdu-guard
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server group grp1 version 2c context context1 vrf management
snmp-server group grp2 version 3 auth context context1 vrf management
snmp-server user user1 grp1 vrf management
snmp-server user user2 grp2 auth md5 encrypt 0x2eaaa9043312c907 vrf management
snmp-server community cx1 vrf management
snmp-server community-map cx1 context context1 user user1 vrf management
feature ntp vrf management
ntp enable vrf management
feature rsyslog vrf management
!
interface ce49
!
```

```
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
!
interface ce54
!
interface ce55
!
interface ce56
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 21.21.21.21/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
ip address 94.94.94.3/24
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6

interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
router ospf 100
```

```
 ospf router-id 21.21.21.21
 snmp context-name context1
 network 21.21.21.21/32 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
!
line console 0
 exec-timeout 0 0
line vty 0 871
 exec-timeout 0 0
 privilege level 16
!
!
end

R1#
```

R2

```
R2(config)#show run
!
! Software version: EC_AS5912-54X-OcNOS-5.1.194-SP-IPBASE-S0-P0 04/20/2022 18:28:5
7
!
!Last configuration change at 15:19:14 UTC Fri Apr 22 2022 by ocnos
!
no service password-encryption
!
logging console 5
logging level all 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
p vrf management
!
hostname R2
no ip domain-lookup
ip domain-lookup vrf management
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context1 vrf management
snmp-server context context2 vrf management
snmp-server group group1 version 3 auth context context1 vrf management
snmp-server group group1 version 3 auth context context2 vrf management
snmp-server user user1 group1 auth md5 encrypt 0x2eaaa9043312c907 vrf management
feature ntp vrf management
ntp enable vrf management
```

```
feature rsyslog vrf management
!
interface ce49
!
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
!
interface ce54
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 23.23.23.23/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
 ip address 94.94.94.2/24
!
interface xe2
 ip address 10.1.2.2/24
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
```

```
router ospf 100
 ospf router-id 23.23.23.23
 snmp context-name context1
 network 23.23.23.23/32 area 0.0.0.0
 network 94.94.94.0/24 area 0.0.0.0
!
router ospf 200
 snmp context-name context2
 network 10.1.2.0/24 area 0.0.0.0
!
end

R2#
```

R3

```
R3#show run
!
! Software version: EC_AS7316-26XB-OcNOS-5.1.194-SP-CSR-S0-P0 04/20/2022 18:28:5
9
!
!Last configuration change at 11:05:18 UTC Mon Feb 25 2019 by ocnos
!
no service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
ip vrf management
!
hostname R3
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server context context2 vrf management
snmp-server group group2 version 3 auth context context2 vrf management
snmp-server group group1 version 2c context context2 vrf management
snmp-server user user2 group2 auth md5 encrypt 0x2eaaa9043312c907 vrf management
snmp-server user user1 group1 vrf management
snmp-server community cx1 vrf management
snmp-server community-map cx1 context context2 user user1 vrf management
feature ntp vrf management
ntp enable vrf management
feature rsyslog vrf management
```

```
!
interface ce0
!
interface ce1
!
interface eth0
 ip vrf forwarding management
 ip address 192.168.3.10/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 24.24.24.24/24 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe0
!
interface xe1
!
interface xe2
 ip address 10.1.2.3/24
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
router ospf 200
 ospf router-id 24.24.24.24
 snmp context-name context2
 network 10.1.2.0/24 area 0.0.0.0
 network 24.24.24.0/24 area 0.0.0.0
!
!
end
```

```
R3#
```

SNMP WALK Command

Perform snmpwalk as mentioned below with IPv4 address using SNMPv3 for R1:

```
snmpwalk -v 3 -u user2 -l auth -r 0 -t 10 -n "context1" -a MD5 -A password  10.12.86.111
iso.3.6.1.2.1
```

Perform snmpwalk as mentioned below with IPv4 address using SNMPv2/SNMPv3 for R2:

```
snmpwalk -v 3 -u user1 -l auth -r 0 -t 10 -n "context2" -a MD5 -A password  10.12.86.116
iso.3.6.1.2.1
```

Perform snmpwalk as mentioned below with IPv4 address using SNMPv3 for R3:

```
snmpwalk -v 3 -u user2 -l auth -r 0 -t 10 -n "context2" -a MD5 -A password  10.12.86.132
iso.3.6.1.2.1
```

OSPF Stub Router Advertisement

The OSPF Stub router feature lets you to prevent a new router in any network from being used as a transit router.

This feature enables the new router to originate link-state advertisement(LSAs) with a maximum metric, thereby causing other routers in the network to prefer alternate-paths instead of using the new router as a transit path to forward traffic. Thereby this feature allows switching and routing functions to be up and running and routing tables to converge before transit traffic is routed again through this new router

Also a startup timer can be configured on this new router enabling it to advertise a maximum metric when the router is started or reloaded

The configuration given below enables the router R3 to advertise router-LSAs/summary-LSAs with maximum metric two routers in Area 0 connecting to network 10.10.10.0/24.

Note: Configure one interface so that it belongs to only one area. It is possible, however, to configure different interfaces on a router to belong to different areas.

Topology

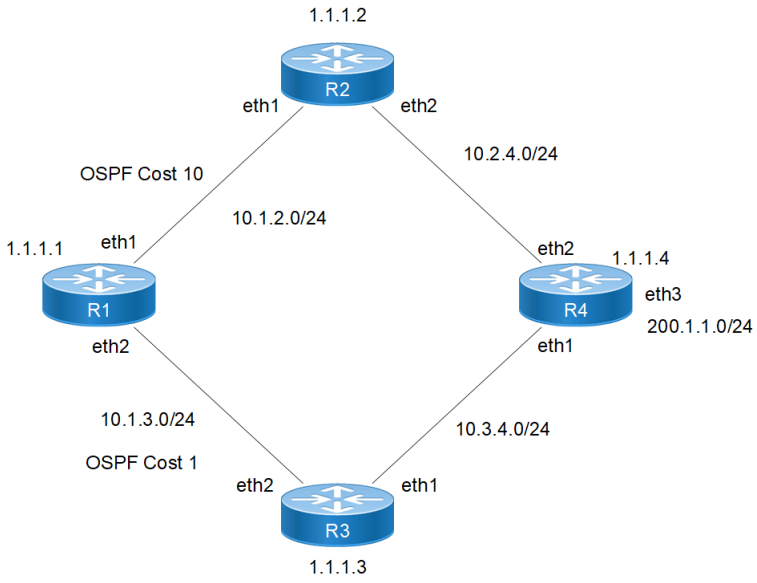

Figure 1-44: OSPF Stub Router Advertisement Topology

Configuration

R1

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config)#ospf router-id 1.1.1.1`	Configure the OSPF router-id
`(config-router)#network 10.1.2.0/24 area 0`	Define the interface (10.1.2.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.1.3.0/24 area 0`	Define the interface (10.1.3.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit router mode
`(config)#interface eth1`	Enter the interface configuration mode for interface eth1
`(config-if)#ip ospf cost 10`	Configure the OSPF cost for the interface
`(config-if)#exit`	Exit the interface configuration mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config)#ospf router-id 1.1.1.2`	Configure the OSPF router-id
`(config-router)#network 10.1.2.0/24 area 0`	Define the interface (10.1.2.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.2.4.0/24 area 1`	Define the interface (10.2.4.0/24) on which OSPF runs, and associate the area ID (1) with the interface
`(config-router)#exit`	Exit router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

R3

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config)#ospf router-id 1.1.1.3`	Configure the OSPF router-id
`(config-router)#network 1.1.1.3/32 area 0`	Define the interface (1.1.1.3/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.1.3.0/24 area 0`	Define the interface (10.1.3.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.3.4.0/24 area 1`	Define the interface (10.3.4.0/24) on which OSPF runs, and associate the area ID (1) with the interface
`(config-router)#max-metric router-lsa include-stub summary-lsa 100`	Configure the router to advertise max-metric for router-LSAs, for stub-networks of router-LSAs and to advertise a metric value of 100 for summary-LSAs
`(config-router)# max-metric router-lsa on-startup 300 include-stub summary-lsa 222`	Configure the router to advertise max-metric for router-LSAs, for stub-networks of router-LSAs and to advertise a metric value of 222 for summary-LSAs on-startup
`(config-router)#exit`	Exit router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

R4

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Configure the routing process, and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process.
`(config)#ospf router-id 1.1.1.4`	Configure the OSPF router-id

(config-router)#network 10.2.4.0/24 area 1	Define the interface (10.2.4.0/24) on which OSPF runs, and associate the area ID (1) with the interface
(config-router)#network 10.3.4.0/24 area 1	Define the interface (10.3.4.0/24) on which OSPF runs, and associate the area ID (1) with the interface
(config-router)#network 200.1.1.0/24 area 1	Define the interface (200.1.1.0/24) on which OSPF runs, and associate the area ID (1) with the interface
(config-router)#exit	Exit router mode
(config)#commit	Commit the candidate configuration to the running configuration.

Validation

R3

```
R3#show running-config ospf
!
router ospf 1
 ospf router-id 1.1.1.3
 max-metric router-lsa include-stub summary-lsa 100
 max-metric router-lsa on-startup 300 include-stub summary-lsa 222
 network 1.1.1.3/32 area 0.0.0.0
 network 10.1.3.0/24 area 0.0.0.0
 network 10.3.4.0/24 area 0.0.0.1
!
R3#
```

```
R3#show ip ospf
 Routing Process "ospf 1" with ID 1.1.1.3
 Process uptime is 1 hour 20 minutes
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Originating router-LSAs with maximum metric
   Condition: always State: active
      Advertise stub links with maximum metric in router-LSAs
      Advertise summary-LSAs with metric 100
   Condition: on start-up for 300 seconds  State: Inactive
      Advertise stub links with maximum metric in router-LSAs
      Advertise summary-LSAs with metric 222
 Supports Graceful Restart
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 SPF schedule delay initial 0 secs 500 msecs
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incoming current DD exchange neighbors 0/64
 Number of outgoing current DD exchange neighbors 0/64
 Initial LSA throttle delay 0 secs 0 msecs
 Minimum hold time for LSA throttle 5 secs 0 msecs
 Maximum wait time for LSA throttle 5 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 0. Checksum 0x000000
```

```
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 0
 External LSA database is unlimited.
 Number of LSA originated 17
 Number of LSA received 98
 IPFRR per-prefix tiebreakers:
    Name                                Index
    Primary path                        20
    Node Protecting                     30
    Interface disjoint                  60
    Broadcast interface disjoint        70
    Downstream path                     90
    Secondary path                      255
 LFA termination hold-on timer : 0 secs 600 msecs
 Number of areas attached to this router: 2
MemPool - struct ospf lsa             : (0-38) | Total (38/100000)
blk_size:168
MemPool - struct rxmt                 : | Total (0/0) blk_size:8
MemPool - OSPF RIB msg4               : | Total (0/0) blk_size:2920
MemPool - OSPF RIB msg4 nh 1     : | Total (0/0) blk_size:24
MemPool - OSPF RIB msg4 nh 2     : | Total (0/0) blk_size:48
MemPool - OSPF RIB msg4 nh 4     : | Total (0/0) blk_size:96
MemPool - OSPF RIB msg4 nh 8     : | Total (0/0) blk_size:192
MemPool - OSPF RIB msg4 nh 16    : | Total (0/0) blk_size:384
MemPool - OSPF RIB msg4 nh 32    : | Total (0/0) blk_size:768
MemPool - OSPF RIB msg4 nh 64    : | Total (0/0) blk_size:1536
MemPool - OSPF RIB msg4 nh 128   : | Total (0/0) blk_size:3072
MemPool - OSPF RIB msg4 nh 255   : | Total (0/0) blk_size:6120
    Area 0.0.0.0 (BACKBONE)
        Number of interfaces in this area is 2(2)
        Number of fully adjacent neighbors in this area is 1
        Area has no authentication
        SPF algorithm last executed 00:03:05.287 ago
        SPF algorithm executed 12 times
        Number of LSA 18. Checksum 0x070234
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 1
        Number of fully adjacent virtual neighbors through this area is 0
        Area has no authentication
        SPF algorithm last executed 00:03:34.838 ago
        SPF algorithm executed 7 times
        Number of LSA 18. Checksum 0x09c738

R3#


R3#show ip ospf database router self-originate

            OSPF Router with ID (1.1.1.3) (Process ID 1 VRF default)

                Router Link States (Area 0.0.0.0)

  LS age: 188
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  Flags: 0x1 : ABR
  LS Type: Router-LSA
```

```
  Link State ID: 1.1.1.3
  Advertising Router: 1.1.1.3
  LS Seq Number: 8000000b
  Checksum: 0x9c4b
  Length: 48
   Number of Links: 2

    Link connected to: Stub Network
     (Link ID) Network/subnet number: 1.1.1.3
     (Link Data) Network Mask: 255.255.255.255
      Number of TOS metrics: 0
       TOS 0 Metric: 65535

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.1.3.1
     (Link Data) Router Interface address: 10.1.3.2
      Number of TOS metrics: 0
       TOS 0 Metric: 65535


                Router Link States (Area 0.0.0.1)

  LS age: 188
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  Flags: 0x1 : ABR
  LS Type: Router-LSA
  Link State ID: 1.1.1.3
  Advertising Router: 1.1.1.3
  LS Seq Number: 8000000c
  Checksum: 0xc531
  Length: 36
   Number of Links: 1

    Link connected to: a Transit Network
     (Link ID) Designated Router address: 10.3.4.2
     (Link Data) Router Interface address: 10.3.4.1
      Number of TOS metrics: 0
       TOS 0 Metric: 65535

R3#


R3#show ip ospf database summary self-originate

            OSPF Router with ID (1.1.1.3) (Process ID 1 VRF default)

                Summary Link States (Area 0.0.0.0)

  LS age: 204
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Summary-LSA
  Link State ID: 10.2.4.0 (summary Network Number)
  Advertising Router: 1.1.1.3
  LS Seq Number: 80000009
  Checksum: 0x7446
  Length: 28
  Network Mask: /24
```

```
      TOS: 0  Metric: 100

LS age: 204
Options: 0x22 (-|-|DC|-|-|-|E|-)
LS Type: Summary-LSA
Link State ID: 10.3.4.0 (summary Network Number)
Advertising Router: 1.1.1.3
LS Seq Number: 80000008
Checksum: 0x6a50
Length: 28
Network Mask: /24
      TOS: 0  Metric: 100

LS age: 204
Options: 0x22 (-|-|DC|-|-|-|E|-)
LS Type: Summary-LSA
Link State ID: 200.1.1.0 (summary Network Number)
Advertising Router: 1.1.1.3
LS Seq Number: 80000009
Checksum: 0xf10e
Length: 28
Network Mask: /24
      TOS: 0  Metric: 100


               Summary Link States (Area 0.0.0.1)

LS age: 204
Options: 0x22 (-|-|DC|-|-|-|E|-)
LS Type: Summary-LSA
Link State ID: 1.1.1.3 (summary Network Number)
Advertising Router: 1.1.1.3
LS Seq Number: 80000008
Checksum: 0xfaca
Length: 28
Network Mask: /32
      TOS: 0  Metric: 100

LS age: 203
Options: 0x22 (-|-|DC|-|-|-|E|-)
LS Type: Summary-LSA
Link State ID: 10.1.2.0 (summary Network Number)
Advertising Router: 1.1.1.3
LS Seq Number: 8000000c
Checksum: 0x902a
Length: 28
Network Mask: /24
      TOS: 0  Metric: 100

LS age: 204
Options: 0x22 (-|-|DC|-|-|-|E|-)
LS Type: Summary-LSA
Link State ID: 10.1.3.0 (summary Network Number)
Advertising Router: 1.1.1.3
LS Seq Number: 80000008
Checksum: 0x8d30
Length: 28
```

```
      Network Mask: /24
            TOS: 0  Metric: 100

R3#
```

CHAPTER 2 Multi-Area Redundant Adjacency Configuration

Overview

The Open Shortest Path First (OSPFv2) is an Interior Gateway Protocol (IGP) that uses a link-state methodology, which calculates the distance among the routers (to determine the link cost) in the network and then distributes the routing information to other routers belonging to a single Autonomous System.

In a link-state routing protocol, each participating router in a network maintains an identical database using the same Autonomous System's topology. Each router's database indicates the local state to other routers. It treats itself as a root and constructs a shortest routing path. This path gives a route for each destination in the Autonomous System.

In this protocol, the entire network is divided into sub-domains. Each sub-domain is referred as 'Area'. One of the sub-domains is classified as ‘Area 0‘ which is the ‘Backbone Area’; rest of the Areas are referred as 'Regular Area” or “Non Zero Area”. The routing among the ‘Regular Area’ is via the ‘Backbone Area’ only even though the cost link is high. A node with one ‘Area 0’ and multiple ‘Regular Area’ is named Area Border Router (ABR).

Enabling one or more ‘Regular areas’ to have Multi-Area over an interface that is part of the ‘Backbone Area’, enhances the routing capability to choose more routing paths as against the regular routing path, irrespective of high link cost. This enhances the routing communications from a ‘Regular Area’ to a ‘Regular Area’ in another ABR via ‘Backbone Area’.

This functionality is enhanced to support the Multi-Area over multiple interfaces of the ‘Backbone Area’.

For more information, refer to https://datatracker.ietf.org/doc/html/rfc5185.

Feature Characteristics

Supports Multi-Area redundancy configuration in multiple interfaces of the ‘Backbone Area’ (aka 'Area 0') for the same 'Regular Area'.

After configuring Multi-Area redundancy, the router provides redundancy paths to reach other networks from the ‘Regular Area’ via ‘Backbone Area’.

The Multi-Area redundancy configuration works only when Multi-Area neighbor is specified using `ip ospf <id> multi-area <area-id> neighbor <neighbor-ip>` command.

LIMITATION:

- It is mandatory to specify the `neighbor ip address` to configure Multi-Area redundancy in more than one interface of ‘Area 0’.
- Only IPv4 support is provided;.
- Configuring more than one Multi-Area in a interface for the same ‘Area 0’ is not supported
- OSPF neighborship should exists to form Multi-Area adjacency
- Multi-Area adjacency is formed between ABR’s.

Benefits

Provides flexibility to use multiple interfaces of Area 0 to reach the Regular Area in different network. Hence, when one of the interfaces goes down the other interface is used to reach the destination.

© 2025 IP Infusion Inc. Proprietary

Prerequisites

The router must be configured with OSPFv2 with Multi-Area support. For more information, refer to Multi-Area Adjacency Configuration.

Configuration

The following configuration enables the Multi-Area feature on multiple interfaces in the OcNOS devices.

Topology

The following scenario illustrates a Regular Area 99 configured between R1-R4, R4-R3, R3-R2 and R2-R5.

The default routing communication between R1 to R3 is using the Multi-Area 99 via R1-R4-R3 even though the link cost between R4-R3 is high.

Figure 2-45 illustrates the Multi-Area Adjacency routing information.

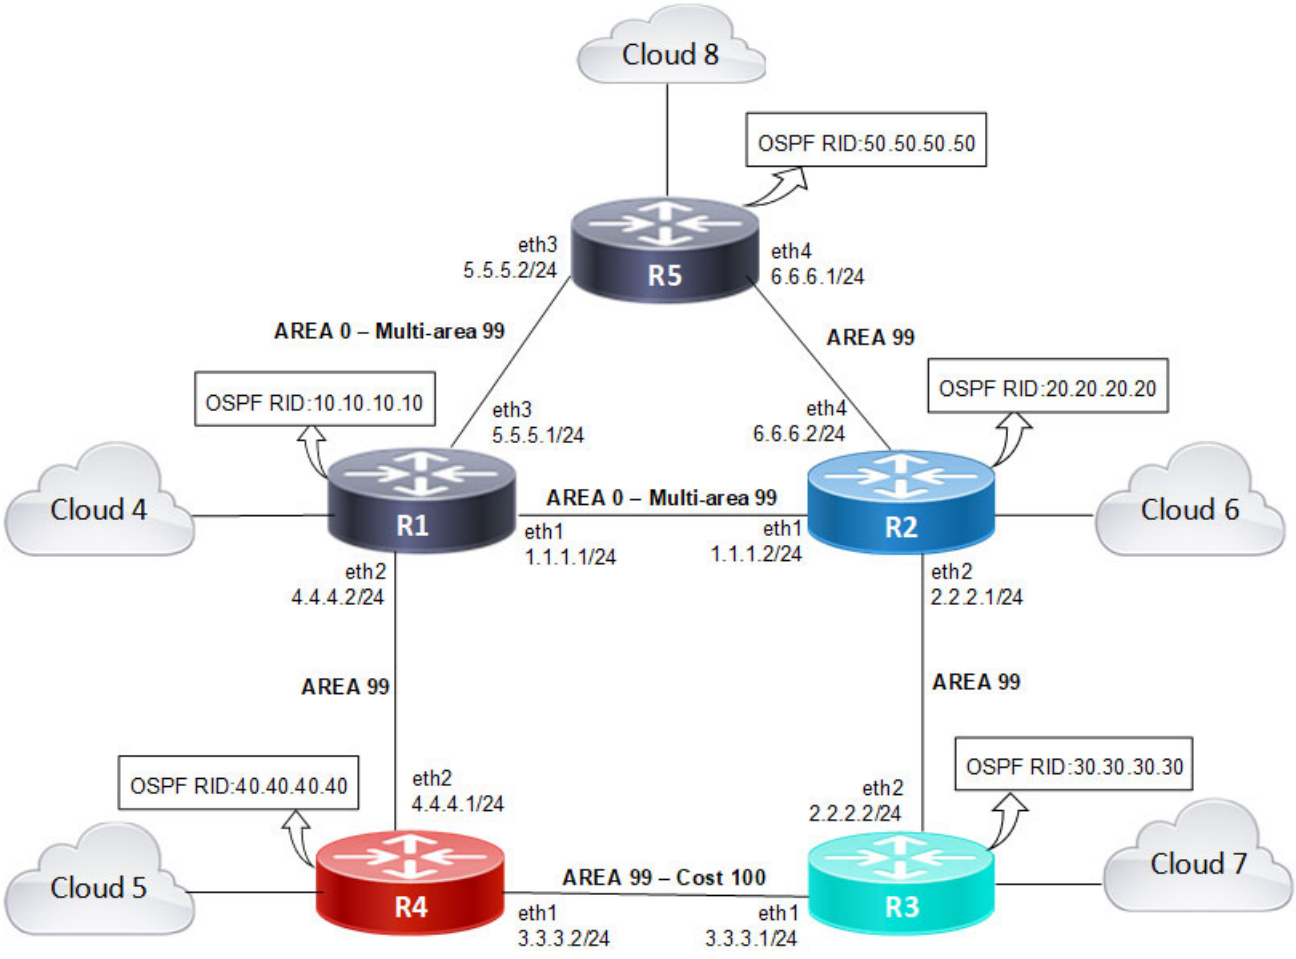

Figure 2-45: OSPFv2 Multi-Area Adjacency with Multiple Interfaces

The Multi-Area support changes this default behavior and permits that once Multi-Area is enabled on `eth1` Area 0 for Multi-Area 99 in R1, the routing communication between R1 and R3 is via R1-R2-R3 as `eth1` Area 99 link cost is high.

Enabling Multi-Area in multiple interfaces, for example on `eth1` and `eth3` of Area 0 with Multi-Area 99, provides redundancy when `eth1` Area 0 goes down. Hence, the routing communication between R1 and R3 occurs via R1-R5-R2-R3.

Perform the following steps to configure the Multi-Area Adjacency on OcNOS devices:

1. Configure the OSPFv2 router's source interfaces IP address, network type, Multi-Area information and OSPF cost of the link on R1, R2, R3, R4, and R5 using the following sample configurations:

R1 source interfaces `eth1, eth2, eth3`

```
(config)#interface eth1
(config-if)#ip address 1.1.1.1/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf 0 multi-area 99

(config-if)#interface eth2
(config-if)#ip address 4.4.4.2/24
(config-if)#ip ospf network point-to-point

(config-if)#interface eth3
(config-if)#ip address 5.5.5.1/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf 0 multi-area 99  neighbor 5.5.5.2
```

R2 source interfaces `eth1, eth2, eth4`

```
(config-if)#interface eth2
(config-if)#ip address 2.2.2.1/24
(config-if)#ip ospf network point-to-point

(config)#interface eth1
(config-if)#ip address 1.1.1.2/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf 0 multi-area 99

(config-if)#interface eth4
(config-if)#ip address 6.6.6.2/24
(config-if)#ip ospf network point-to-point
```

R3 source interfaces `eth1, eth2`

```
(config)#interface eth1
(config-if)#ip address 3.3.3.1/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf cost 100

(config)#interface eth2
(config-if)#ip address 2.2.2.2/24
(config-if)#ip ospf network point-to-point
```

R4 source interfaces `eth1, eth2`

```
(config)#interface eth1
(config-if)#ip address 3.3.3.2/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf cost
(config)#interface eth2
(config-if)#ip address 4.4.4.1/24
(config-if)#ip ospf network point-to-point
```

R5 source interfaces `eth3, eth4`

```
(config)#interface eth3
(config-if)#ip address 5.5.5.2/24
(config-if)#ip ospf network point-to-point
(config-if)#ip ospf 0 multi-area 0.0.0.99

(config-if)#interface eth4
(config-if)#ip address 6.6.6.1/24
(config-if)#ip ospf network point-to-point
```

2. Configure an OSPF instance with an instance ID on R1, R2, R3, R4, and R5. Refer to theTopology diagram for the router ID.

```
(config)#router ospf
(config-router)# ospf router-id 10.10.10.10
```

3. Configure OSPF routing with Multi-Area area 0 /area 99 area ID and assign loopback IP address to R1, R2, R3, R4 and R5.

R1 - between R1 - R2, R1 - R4, and R1 - R5

```
(config-router)#network 1.1.1.0/24 area 0.0.0.0
(config-router)#network 4.4.4.0/24 area 0.0.0.99
(config-router)#network 5.5.5.0/24 area 0.0.0.0
(config-router)#network 10.10.10.10/32 area 0.0.0.0
```

R2 - between R1 - R2, R2 - R3, and R2 - R4

```
(config-router)#network 1.1.1.0/24 area 0.0.0.0
(config-router)#network 2.2.2.0/24 area 0.0.0.99
(config-router)#network 6.6.6.0/24 area 0.0.0.99
(config-router)#network 20.20.20.20/32 area 0.0.0.99
```

R3 - between R2 - R3, and R3 - R4

```
(config-router)#network 2.2.2.0/24 0.0.0.99
(config-router)#network 3.3.3.0/24 area 0.0.0.99
(config-router)#network 30.30.30.30/32 area 0.0.0.99
```

R4 - between R3 - R4, and R1 - R4.

```
(config-router)#network 3.3.3.0/24 area 0.0.0.99
(config-router)#network 4.4.4.0/24 area 0.0.0.99
(config-router)#network 40.40.40.40/32 area 0.0.0.99
```

R5 - between R1 - R5, and R2 - R5.

```
(config-router)#network 5.5.5.0/24 area 0.0.0.0
(config-router)#network 6.6.6.0/24 area 0.0.0.99
(config-router)#network 50.50.50.50/32 area 0.0.0.99
```

Sample show running-config Output

R1

```
!
interface eth1
 ip address 1.1.1.1/24
 ip ospf network point-to-point
 ip ospf 0 multi-area 0.0.0.99
!
interface eth2
 ip address 4.4.4.2/24
 ip ospf network point-to-point
!
interface eth3
 ip address 5.5.5.1/24
 ip ospf network point-to-point
 ip ospf 0 multi-area 0.0.0.99 neighbor 5.5.5.2
!


!
router ospf
 ospf router-id 10.10.10.10
 network 1.1.1.0/24 area 0.0.0.0
 network 4.4.4.0/24 area 0.0.0.99
 network 5.5.5.0/24 area 0.0.0.0
 network 10.10.10.10/32 area 0.0.0.0
!
!
```

R2

```
!
interface eth1
 ip address 1.1.1.2/24
 ip ospf network point-to-point
 ip ospf 0 multi-area 0.0.0.99
!
interface eth2
 ip address 2.2.2.1/24
 ip ospf network point-to-point
!
interface eth3
!
interface eth4
 ip address 6.6.6.2/24
 ip ospf network point-to-point
!
router ospf
 ospf router-id 20.20.20.20
```

```
 network 1.1.1.0/24 area 0.0.0.0
 network 2.2.2.0/24 area 0.0.0.99
 network 6.6.6.0/24 area 0.0.0.99
 network 20.20.20.20/32 area 0.0.0.99
!
```

R3

```
!
interface eth1
 ip address 3.3.3.1/24
 ip ospf network point-to-point
 ip ospf cost 250
!
interface eth2
 ip address 2.2.2.2/24
 ip ospf network point-to-point
!
!
router ospf
 ospf router-id 30.30.30.30
 network 2.2.2.0/24 area 0.0.0.99
 network 3.3.3.0/24 area 0.0.0.99
 network 30.30.30.30/32 area 0.0.0.99
!
```

R4

```
!
interface eth1
 ip address 3.3.3.2/24
 ip ospf network point-to-point
 ip ospf cost 250
!
interface eth2
 ip address 4.4.4.1/24
 ip ospf network point-to-point
!
router ospf
 ospf router-id 40.40.40.40
 network 3.3.3.0/24 area 0.0.0.99
 network 4.4.4.0/24 area 0.0.0.99
 network 40.40.40.40/32 area 0.0.0.99
!
ip route 20.1.1.0/24 10.1.1.1
!
!
```

R5

```
!
interface eth3
```

```
 ip address 5.5.5.2/24
 ip ospf network point-to-point
 ip ospf 0 multi-area 0.0.0.99
!
interface eth4
 ip address 6.6.6.1/24
 ip ospf network point-to-point
!
!
router ospf
 ospf router-id 50.50.50.50
 network 5.5.5.0/24 area 0.0.0.0
 network 6.6.6.0/24 area 0.0.0.99
 network 50.50.50.50/32 area 0.0.0.99
!
```

Validation

R1

The following show output displays the Multi-Area Adjacency information for OSPF router.

```
OcNOS#show ip ospf multi-area-adjacencies
  Multi-area-adjacency link on interface eth1 to neighbor 1.1.1.2
  Internet Address 1.1.1.1/24, Area 0.0.0.99, MTU 1500
  Process ID 0, Router ID 10.10.10.10, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:01
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 700 sent 743, DD received 3 sent 4
  LS-Req received 0 sent 0, LS-Upd received 53 sent 80
  LS-Ack received 50 sent 25, Discarded 0
  Multi-area-adjacency link on interface eth3 to neighbor 5.5.5.2
  Internet Address 5.5.5.1/24, Area 0.0.0.99, MTU 1500
  Process ID 0, Router ID 10.10.10.10, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:07
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 684 sent 735, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 31 sent 79
  LS-Ack received 70 sent 23, Discarded 0
OcNOS#
```

The following show output displays the neighbor IP address of the R1.

```
OcNOS#show ip ospf neighbor
      Total number of full neighbors: 5
      OSPF process 0 VRF(default):
```

```
      Neighbor ID     Pri   State    Dead Time   Address      Interface      Instance ID
      20.20.20.20       1   Full/ -   00:00:32    1.1.1        eth1           0
      40.40.40.40       1   Full/ -   00:00:29    4.4.4.1      eth2           0
      50.50.50.50       1   Full/ -   00:00:39    5.5.5.2      eth3           0
      20.20.20.20       1   Full/ -   00:00:36    1.1.1.2      eth1
      50.50.50.50       1   Full/ -   00:00:37    5.5.5.2      eth3
OcNOS#
```

The following show output displays the OSPF routing information of the R1.

```
OcNOS#show ip ospf route

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

C  1.1.1.0/24 [1] is directly connected, eth1, Area 0.0.0.0
O  2.2.2.0/24 [2] via 1.1.1.2, eth1, Area 0.0.0.99
O  3.3.3.0/24 [101] via 4.4.4.1, eth2, Area 0.0.0.99
C  4.4.4.0/24 [1] is directly connected, eth2, Area 0.0.0.99
C  5.5.5.0/24 [1] is directly connected, eth3, Area 0.0.0.0
O  6.6.6.0/24 [2] via 1.1.1.2, eth1, Area 0.0.0.99
                  via 5.5.5.2, eth3, Area 0.0.0.99
C  10.10.10.10/32 [1] is directly connected, lo, Area 0.0.0.0
O  20.20.20.20/32 [2] via 1.1.1.2, eth1, Area 0.0.0.99
O  30.30.30.30/32 [3] via 1.1.1.2, eth1, Area 0.0.0.99
O  40.40.40.40/32 [2] via 4.4.4.1, eth2, Area 0.0.0.99
O  50.50.50.50/32 [2] via 5.5.5.2, eth3, Area 0.0.0.99
OcNOS#
```

The following show output displays the IP route information of the R1.

```
      OcNOS#show ip route
      Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
             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
             i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
             ia - IS-IS inter area, E - EVPN,
             v - vrf leaked
             * - candidate default

      IP Route Table for VRF "default"
      Gateway of last resort is 10.12.100.1 to network 0.0.0.0

      K*           0.0.0.0/0 [0/0] via 10.12.100.1, eth0
```

```
C            1.1.1.0/24 is directly connected, eth1, 01:54:16
O            2.2.2.0/24 [110/2] via 1.1.1.2, eth1, 01:23:42
O            3.3.3.0/24 [110/101] via 4.4.4.1, eth2, 00:11:12
C            4.4.4.0/24 is directly connected, eth2, 5d09h32m
C            5.5.5.0/24 is directly connected, eth3, 5d09h39m
O            6.6.6.0/24 [110/2] via 5.5.5.2, eth3, 01:23:57
                        [110/2] via 1.1.1.2, eth1
C            10.10.10.10/32 is directly connected, lo, 00:32:27
C            10.12.100.0/22 is directly connected, eth0, 03w1d02h
O            20.20.20.20/32 [110/2] via 1.1.1.2, eth1, 00:27:24
O            30.30.30.30/32 [110/3] via 1.1.1.2, eth1, 00:21:38
O            40.40.40.40/32 [110/2] via 4.4.4.1, eth2, 00:17:09
O            50.50.50.50/32 [110/2] via 5.5.5.2, eth3, 00:04:53
C            127.0.0.0/8 is directly connected, lo, 03w1d02h
OcNOS#
```

R2

The following show output displays the Multi-Area Adjacency information for OSPF router.

```
OcNOS#show ip ospf multi-area-adjacencies
  Multi-area-adjacency link on interface eth1 to neighbor 1.1.1.1
  Internet Address 1.1.1.2/24, Area 0.0.0.99, MTU 1500
  Process ID 0, Router ID 20.20.20.20, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:00
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 726 sent 727, DD received 4 sent 3
  LS-Req received 0 sent 0, LS-Upd received 85 sent 56
  LS-Ack received 25 sent 52, Discarded 0
OcNOS#
```

The following show output displays the neighbor IP address of the R2.

```
OcNOS#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 0 VRF(default):
Neighbor ID     Pri   State        Dead Time   Address        Interface       Instance ID
10.10.10.10       1   Full/ -      00:00:33    1.1.1.1          eth1           0
30.30.30.30       1   Full/ -      00:00:29    2.2.2.2          eth2           0
50.50.50.50       1   Full/ -      00:00:32    6.6.6.1          eth4           0
10.10.10.10       1   Full/ -      00:00:35    1.1.1.1          eth1           0
OcNOS#
```

The following show output displays the OSPF Multi-Area routing information of the R1.

```
OcNOS#show ip ospf route

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

C  1.1.1.0/24 [1] is directly connected, eth1, Area 0.0.0.0
C  2.2.2.0/24 [1] is directly connected, eth2, Area 0.0.0.99
O  3.3.3.0/24 [101] via 2.2.2.2, eth2, Area 0.0.0.99
O  4.4.4.0/24 [2] via 1.1.1.1, eth1, Area 0.0.0.99
O  5.5.5.0/24 [2] via 1.1.1.1, eth1, Area 0.0.0.0
C  6.6.6.0/24 [1] is directly connected, eth4, Area 0.0.0.99
O  10.10.10.10/32 [2] via 1.1.1.1, eth1, Area 0.0.0.0
C  20.20.20.20/32 [1] is directly connected, lo, Area 0.0.0.99
O  30.30.30.30/32 [2] via 2.2.2.2, eth2, Area 0.0.0.99
O  40.40.40.40/32 [3] via 1.1.1.1, eth1, Area 0.0.0.99
O  50.50.50.50/32 [2] via 6.6.6.1, eth4, Area 0.0.0.99
```

The following show output displays the IP routing information of the R2.

```
OcNOS#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.100.1 to network 0.0.0.0

K*           0.0.0.0/0 [0/0] via 10.12.100.1, eth0
C            1.1.1.0/24 is directly connected, eth1, 5d09h39m
C            2.2.2.0/24 is directly connected, eth2, 5d09h35m
O            3.3.3.0/24 [110/101] via 2.2.2.2, eth2, 00:20:15
O            4.4.4.0/24 [110/2] via 1.1.1.1, eth1, 01:36:00
O            5.5.5.0/24 [110/2] via 1.1.1.1, eth1, 01:36:05
C            6.6.6.0/24 is directly connected, eth4, 5d09h37m
O            10.10.10.10/32 [110/2] via 1.1.1.1, eth1, 00:44:39
C            10.12.100.0/22 is directly connected, eth0, 03w1d02h
C            20.20.20.20/32 is directly connected, lo, 00:39:37
O            30.30.30.30/32 [110/2] via 2.2.2.2, eth2, 00:33:51
O            40.40.40.40/32 [110/3] via 1.1.1.1, eth1, 00:29:22
O            50.50.50.50/32 [110/2] via 6.6.6.1, eth4, 00:17:06
C            127.0.0.0/8 is directly connected, lo, 03w1d02h
OcNOS#
```

R3

The following show output displays the Multi-Area Adjacency information for OSPF R3 router.

```
OcNOS#show ip ospf multi-area-adjacencies
OcNOS#

OcNOS#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 0 VRF(default):
Neighbor ID     Pri   State    Dead Time   Address     Interface     Instance ID
20.20.20.20       1   Full/ -   00:00:29    2.2.2.1     eth2          0
40.40.40.40       1   Full/ -   00:00:35    3.3.3.2     eth1          0
OcNOS#
```

The following show output displays the OSPF routing information of the R3.

```
OcNOS#show ip ospf route

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

IA 1.1.1.0/24 [2] via 2.2.2.1, eth2, Area 0.0.0.99
C  2.2.2.0/24 [1] is directly connected, eth2, Area 0.0.0.99
C  3.3.3.0/24 [100] is directly connected, eth1, Area 0.0.0.99
O  4.4.4.0/24 [3] via 2.2.2.1, eth2, Area 0.0.0.99
IA 5.5.5.0/24 [3] via 2.2.2.1, eth2, Area 0.0.0.99
O  6.6.6.0/24 [2] via 2.2.2.1, eth2, Area 0.0.0.99
IA 10.10.10.10/32 [3] via 2.2.2.1, eth2, Area 0.0.0.99
O  20.20.20.20/32 [2] via 2.2.2.1, eth2, Area 0.0.0.99
C  30.30.30.30/32 [1] is directly connected, lo, Area 0.0.0.99
O  40.40.40.40/32 [4] via 2.2.2.1, eth2, Area 0.0.0.99
O  50.50.50.50/32 [3] via 2.2.2.1, eth2, Area 0.0.0.99
OcNOS#
```

The following show output displays the IP routing information of the R3.

```
OcNOS#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
```

```
O IA         1.1.1.0/24 [110/2] via 4.4.4.2, eth2, 02:10:02
O            2.2.2.0/24 [110/3] via 4.4.4.2, eth2, 02:09:57
C            3.3.3.0/24 is directly connected, eth1, 5d09h36m
C            4.4.4.0/24 is directly connected, eth2, 5d09h36m
O IA         5.5.5.0/24 [110/2] via 4.4.4.2, eth2, 02:10:02
O            6.6.6.0/24 [110/3] via 4.4.4.2, eth2, 02:10:02
O IA         10.10.10.10/32 [110/2] via 4.4.4.2, eth2, 01:18:36
C            10.12.100.0/22 is directly connected, eth0, 05w1d18h
O            20.20.20.20/32 [110/3] via 4.4.4.2, eth2, 01:13:33
O            30.30.30.30/32 [110/4] via 4.4.4.2, eth2, 01:07:47
C            40.40.40.40/32 is directly connected, lo, 01:03:19
O            50.50.50.50/32 [110/3] via 4.4.4.2, eth2, 00:51:02
C            127.0.0.0/8 is directly connected, lo, 05w1d18h
```

R4

The following show output displays the OSPF routing information of the R4.

```
OcNOS#show ip ospf route

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

IA 1.1.1.0/24 [2] via 4.4.4.2, eth2, Area 0.0.0.99
O  2.2.2.0/24 [3] via 4.4.4.2, eth2, Area 0.0.0.99
C  3.3.3.0/24 [100] is directly connected, eth1, Area 0.0.0.99
C  4.4.4.0/24 [1] is directly connected, eth2, Area 0.0.0.99
IA 5.5.5.0/24 [2] via 4.4.4.2, eth2, Area 0.0.0.99
O  6.6.6.0/24 [3] via 4.4.4.2, eth2, Area 0.0.0.99
IA 10.10.10.10/32 [2] via 4.4.4.2, eth2, Area 0.0.0.99
O  20.20.20.20/32 [3] via 4.4.4.2, eth2, Area 0.0.0.99
O  30.30.30.30/32 [4] via 4.4.4.2, eth2, Area 0.0.0.99
C  40.40.40.40/32 [1] is directly connected, lo, Area 0.0.0.99
O  50.50.50.50/32 [3] via 4.4.4.2, eth2, Area 0.0.0.99
OcNOS#
```

The following show output displays the IP routing information of the R4.

```
OcNOS#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
```

```
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O IA         1.1.1.0/24 [110/2] via 4.4.4.2, eth2, 02:10:02
O            2.2.2.0/24 [110/3] via 4.4.4.2, eth2, 02:09:57
C            3.3.3.0/24 is directly connected, eth1, 5d09h36m
C            4.4.4.0/24 is directly connected, eth2, 5d09h36m
O IA         5.5.5.0/24 [110/2] via 4.4.4.2, eth2, 02:10:02
O            6.6.6.0/24 [110/3] via 4.4.4.2, eth2, 02:10:02
O IA         10.10.10.10/32 [110/2] via 4.4.4.2, eth2, 01:18:36
C            10.12.100.0/22 is directly connected, eth0, 05w1d18h
O            20.20.20.20/32 [110/3] via 4.4.4.2, eth2, 01:13:33
O            30.30.30.30/32 [110/4] via 4.4.4.2, eth2, 01:07:47
C            40.40.40.40/32 is directly connected, lo, 01:03:19
O            50.50.50.50/32 [110/3] via 4.4.4.2, eth2, 00:51:02
C            127.0.0.0/8 is directly connected, lo, 05w1d18h

Gateway of last resort is not set

OcNOS#
```

R5

The following show output displays the Multi-Area Adjacency information for OSPF router.

```
OcNOS#show ip ospf multi-area-adjacencies
  Multi-area-adjacency link on interface eth3 to neighbor 5.5.5.1
  Internet Address 5.5.5.2/24, Area 0.0.0.99, MTU 1500
  Process ID 0, Router ID 50.50.50.50, Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:08
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 779 sent 782, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 87 sent 41
  LS-Ack received 30 sent 75, Discarded 0
OcNOS#
```

The following show output displays the neighbor IP address of the R2.

```
OcNOS#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 0 VRF(default):
Neighbor ID     Pri   State  Dead Time   Address      Interface     Instance ID
10.10.10.10      1    Full/ - 00:00:35    5.5.5.1     eth3          0
20.20.20.20      1    Full/  - 00:00:34    6.6.6.2     eth4          0
10.10.10.10      1    Full/ - 00:00:32    5.5.5.1     eth3          0
```

```
OcNOS#
```

The following show output displays the OSPF routing information of the R5.

```
OcNOS#show ip ospf route

OSPF process 0:
Codes: C - connected, D - Discard, 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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

O  1.1.1.0/24 [2] via 5.5.5.1, eth3, Area 0.0.0.0
O  2.2.2.0/24 [2] via 6.6.6.2, eth4, Area 0.0.0.99
O  3.3.3.0/24 [102] via 6.6.6.2, eth4, Area 0.0.0.99
                    via 5.5.5.1, eth3, Area 0.0.0.99
O  4.4.4.0/24 [2] via 5.5.5.1, eth3, Area 0.0.0.99
C  5.5.5.0/24 [1] is directly connected, eth3, Area 0.0.0.0
C  6.6.6.0/24 [1] is directly connected, eth4, Area 0.0.0.99
O  10.10.10.10/32 [2] via 5.5.5.1, eth3, Area 0.0.0.0
O  20.20.20.20/32 [2] via 6.6.6.2, eth4, Area 0.0.0.99
O  30.30.30.30/32 [3] via 6.6.6.2, eth4, Area 0.0.0.99
O  40.40.40.40/32 [3] via 5.5.5.1, eth3, Area 0.0.0.99
C  50.50.50.50/32 [1] is directly connected, lo, Area 0.0.0.99
OcNOS#
```

The following show output displays the IP routing information of the R5.

```
OcNOS#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            1.1.1.0/24 [110/2] via 5.5.5.1, eth3, 01:47:39
O            2.2.2.0/24 [110/2] via 6.6.6.2, eth4, 5d03h27m
O            3.3.3.0/24 [110/102] via 5.5.5.1, eth3, 5d04h59m
                        [110/102] via 6.6.6.2, eth4
O            4.4.4.0/24 [110/2] via 5.5.5.1, eth3, 01:47:39
C            5.5.5.0/24 is directly connected, eth3, 5d09h36m
C            6.6.6.0/24 is directly connected, eth4, 5d09h34m
O            10.10.10.10/32 [110/2] via 5.5.5.1, eth3, 00:56:06
```

```
C            10.12.100.0/22 is directly connected, eth0, 04w0d23h
O            20.20.20.20/32 [110/2] via 6.6.6.2, eth4, 00:51:04
O            30.30.30.30/32 [110/3] via 6.6.6.2, eth4, 00:45:18
O            40.40.40.40/32 [110/3] via 5.5.5.1, eth3, 00:40:49
C            50.50.50.50/32 is directly connected, lo, 00:28:33
C            127.0.0.0/8 is directly connected, lo, 04w0d23h

Gateway of last resort is not set
OcNOS#
```

Implementation Example

The OSPF Multi-area is used in high-speed link between two Area Border Routers (ABRs) in multiple areas.

Revised CLI Commands

The following existing CLI is updated to support this feature.

ip ospf multi-area

Use this command to enable Multi-Area adjacency on point-to-point network and other network types. Multi-Area adjacency establishes adjacency between the Area Border Routers (ABRs). The interface of the ABR where this command is configured, shall be associated with multiple areas.

Use the `no` parameter to disable Multi-Area adjacency on the given interface on point-to-point network.

Command Syntax

```
ip ospf <0-65535> multi-area (A.B.C.D|<0-4294967295>)(neighbor A.B.C.D |)
no ip ospf <0-65535> multi-area (A.B.C.D|<0-4294967295>) (neighbor|)
```

Parameters

`ospf <0-65535>`	OSPF process ID.
`multi-area`	
`A.B.C.D`	OSPF area ID in IP address format.
`<0-4294967295>`	OSPF area ID as a decimal value.
`neighbor`	
`A.B.C.D`	Neighbor IP address.

Note: The neighbor parameter is mandatory to create redundant Multi-Area adjacencies in one Area.

Default

The Multi-Area adjacency is disabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3. In OcNOS version 6.5.1, the neighbor parameter is made mandatory to create redundant Multi-Area adjacencies in one Area.

Examples

```
#configure terminal
(config)#interface eth1
(config-if)#ip ospf 0 multi-area 1
(config-if)# no ip ospf 0 multi-area 1
```

To configure redundant Multi-Area adjacency configuration with same area in an OSPF instance.

```
#configure terminal
(config)#interface eth1
(config-if)#ip ospf 0 multi-area 1
(config)#interface eth2
(config-if)#ip ospf 0 multi-area 1 neighbor 5.5.5.2
(config)#interface eth3
(config-if)#ip ospf 0 multi-area 1 neighbor 3.3.3.2
(config-if)# no ip ospf 0 multi-area 1
```

Glossary

Key Terms/ Acronym	Description
OSPF	An Interior Gateway Protocol (IGP) based on link-state routing. Open Shortest Path First (OSPF) is widely deployed in large networks because of its efficient use of network bandwidth and its rapid convergence after changes in topology. Defined in RFCs 2328 and RFC 5340. OSPF advertises the states of local network links within an autonomous system (AS) and makes routing decisions based on the shortest path first (SPF) algorithm. Each OSPF router maintains an identical database describing the autonomous system's topology. From this database, a Routing Information Base (RIB) is calculated by constructing a shortest path tree (SPT).

IGP	An intradomain protocol used to exchange network reachability and routing information among devices within an autonomous system (AS), such as Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF), or Routing Information Protocol (RIP). Contrast with Exterior Gateway Protocol (EGP)
ABR	A router on the border of one or more Open Shortest Path First (OSPF) areas that connects those areas to the backbone network. An ABR is a member of both the OSPF backbone and its attached areas. Therefore, an ABR maintains routing tables for both the backbone topology and the topology of the other areas. See also Not-So-Stubby-Area (NSSA), stub area.

© 2025 IP Infusion Inc. Proprietary

CHAPTER 3 OSPFv3

This chapter contains basic OSPFv3 configuration examples.

Enable OSPFv3 on an Interface

This example shows the minimum configuration required for enabling OSPFv3 on an interface. R1 and R2 are two routers in Area 0 connecting to the network 3ffe:10::/64. After enabling OSPFv3 on an interface, create a routing instance, and specify the Router ID.

Note: You must explicitly specify a Router ID for the OSPFv3 process to be activated.

Topology

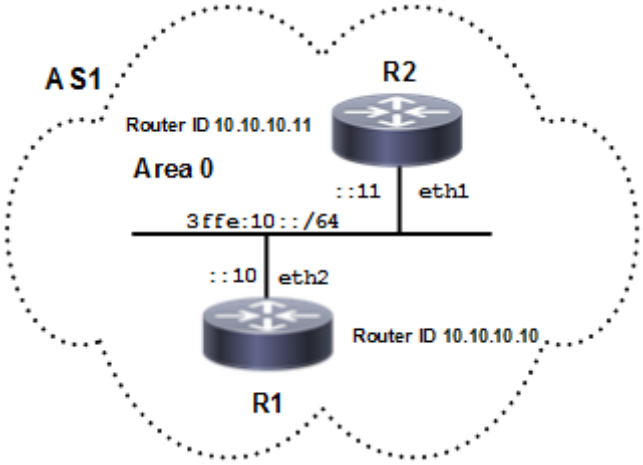

Figure 3-46: Basic OSPFv3 Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.10`	Specify a Router ID for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `0`.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.11`	Specify a Router ID (`10.10.10.11`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.11       1   Full/Backup     00:00:35    eth2       0

#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.10) (Process *null*)

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age        Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10      164       0x80000001 0xf3c6      1
0.0.0.3         10.10.10.11      106       0x80000001 0xd973      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum    Link
0.0.0.0         10.10.10.10       94       0x80000003 0xb2f0      1
0.0.0.0         10.10.10.11       95       0x80000003 0x9e05      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.4         10.10.10.10       94       0x80000001 0xf990

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.10       93       0x80000001 0xc35d      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)
```

```
Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10       94       0x80000002 0x3504
0.0.0.3         10.10.10.11       95       0x80000002 0x6bcc

#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10            --
10.10.10.11            1          10.10.10.11          eth2
```

R2

```
#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/DR         00:00:31    eth1       0

R2#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.11) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10      341       0x80000001 0xf3c6      1
0.0.0.3         10.10.10.11      281       0x80000001 0xd973      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10      271       0x80000003 0xb2f0      1
0.0.0.0         10.10.10.11      270       0x80000003 0x9e05      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10      271       0x80000001 0xf990

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.10      270       0x80000001 0xc35d      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10      271       0x80000002 0x3504
0.0.0.3         10.10.10.11      270       0x80000002 0x6bcc
```

```
R2#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10            1          10.10.10.10          eth1
10.10.10.11            --
```

Set Priority

This example shows how to set priority for an interface. Set a high priority for a router to make it the Designated Router (DR). Router R3 is configured with a priority of 10; this is higher than the default priority (default priority is 1) set for R1 and R2. This makes R3 the DR.

Topology

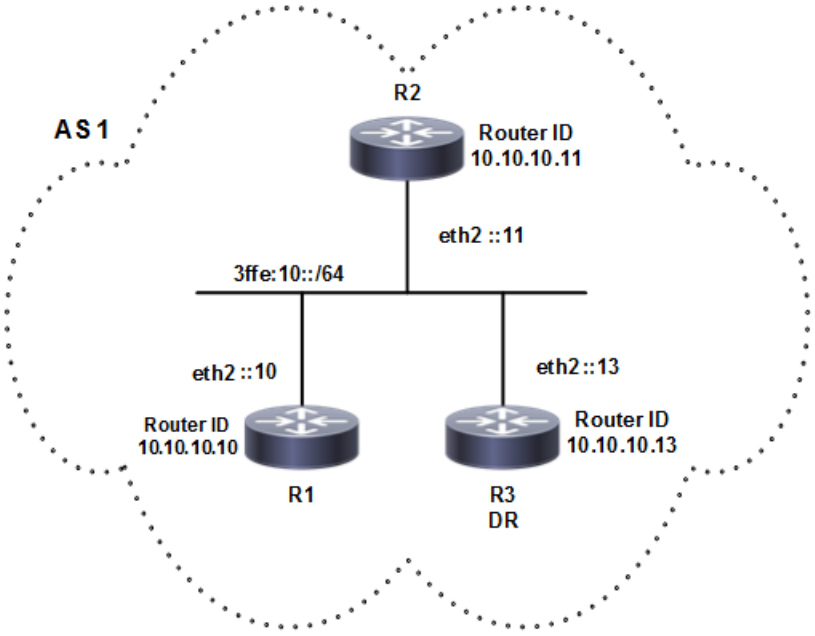

Figure 3-47: OSPFv3 Set Priority

R3

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.13`	Specify a Router ID (`10.10.10.13`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.

`(config-if)#ipv6 ospf priority 10`	Specify the router priority to a higher priority (10) to make R3 the Designated Router (DR).
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.10`	Specify a Router ID (`10.10.10.10`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R2

`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.11`	Specify a Router ID (`10.10.10.11`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
rtr1#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface Instance ID
10.10.10.11       1   Full/DROther    00:00:37    eth2       0
10.10.10.13      10   Full/DR         00:00:37    eth2       0

rtr1#show ipv6 ospf database
            OSPFv3 Router with ID (10.10.10.10) (Process *null*)

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum Prefix
0.0.0.4         10.10.10.10      398       0x80000001 0xf3c6      1
0.0.0.4         10.10.10.11       71       0x80000001 0x4768      1
0.0.0.4         10.10.10.13      611       0x80000002 0x695b      1
```

```
                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10       49       0x80000004 0xf2ac      1
0.0.0.0         10.10.10.11       50       0x80000004 0xecb1      1
0.0.0.0         10.10.10.13       61       0x80000004 0xe0bb      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.13       61       0x80000002 0xa6b0

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.13       60       0x80000002 0xd940      1  Network-
LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10       49       0x80000003 0x75bf
0.0.0.4         10.10.10.11       50       0x80000004 0x9f92
0.0.0.4         10.10.10.13       61       0x80000003 0xf935


rtr1#show ipv6 ospfv3 topology
OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits Metric     Next-Hop             Interface
10.10.10.10            --
10.10.10.11             1          10.10.10.11          eth2
10.10.10.13             1          10.10.10.13          eth2
```

R2

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface Instance ID
10.10.10.10       1   Full/Backup     00:00:31    eth2       0
10.10.10.13      10   Full/DR         00:00:39    eth2       0

R2#show ipv6 ospf database
            OSPFv3 Router with ID (10.10.10.11) (Process *null*)

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10      525       0x80000001 0xf3c6      1
0.0.0.4         10.10.10.11      194       0x80000001 0x4768      1
0.0.0.4         10.10.10.13      736       0x80000002 0x695b      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
```

```
0.0.0.0         10.10.10.10      175       0x80000004 0xf2ac      1
0.0.0.0         10.10.10.11      174       0x80000004 0xecb1      1
0.0.0.0         10.10.10.13      186       0x80000004 0xe0bb      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.13      186       0x80000002 0xa6b0

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.13      185       0x80000002 0xd940      1  Network-
LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10      175       0x80000003 0x75bf
0.0.0.4         10.10.10.11      174       0x80000004 0x9f92
0.0.0.4         10.10.10.13      186       0x80000003 0xf935

R2#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits Metric     Next-Hop             Interface
10.10.10.10            1          10.10.10.10          eth2
10.10.10.11            --
10.10.10.13            1          10.10.10.13          eth2
```

R3

```
R3#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/Backup     00:00:38    eth2       0
10.10.10.11       1   Full/DROther    00:00:29    eth2       0

R3#show ipv6 ospf database
            OSPFv3 Router with ID (10.10.10.13) (Process *null*)

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10      658       0x80000001 0xf3c6      1
0.0.0.4         10.10.10.11      329       0x80000001 0x4768      1
0.0.0.4         10.10.10.13      869       0x80000002 0x695b      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10      309       0x80000004 0xf2ac      1
0.0.0.0         10.10.10.11      309       0x80000004 0xecb1      1
0.0.0.0         10.10.10.13      319       0x80000004 0xe0bb      1

                Network-LSA (Area 0.0.0.0)
```

```
Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.13      319       0x80000002 0xa6b0

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.13      318       0x80000002 0xd940      1  Network-
LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10      309       0x80000003 0x75bf
0.0.0.4         10.10.10.11      309       0x80000004 0x9f92
0.0.0.4         10.10.10.13      319       0x80000003 0xf935

R3#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits Metric     Next-Hop             Interface
10.10.10.10            1          10.10.10.10          eth2
10.10.10.11            1          10.10.10.11          eth2
10.10.10.13            --
```

Area Border Router

This example shows configuration for an Area Border Router. R2 is an Area Border Router (ABR). On R2, interface eth2 is in Area 0, and interface eth1 is in Area 1.

Topology

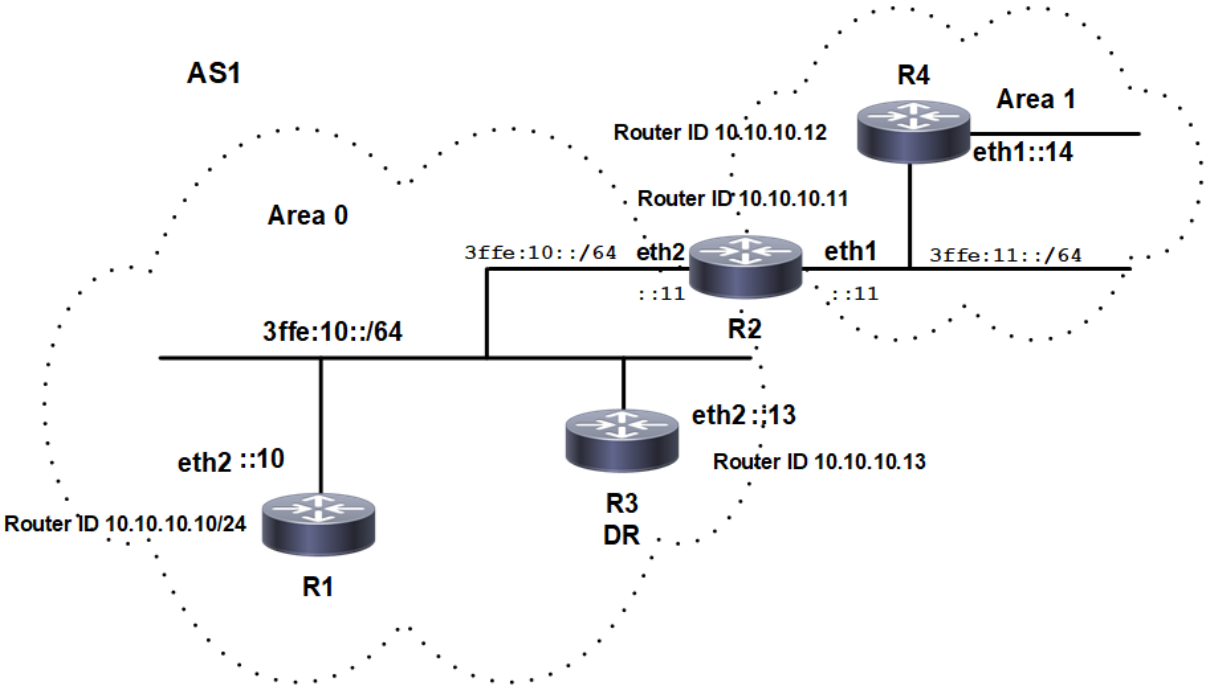

Figure 3-48: OSPFv3 Area Border Router

Configuration

R2

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.11`	Specify a Router ID (`10.10.10.11`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on the other interface, and assign the other Area ID `(1)`.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.

R4

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.12`	Specify a Router ID (10.10.10.12) for the OSPFv3 routing process.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on the other interface, and assign the other Area ID (1).
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.

Validation

R2

```
R2#show ipv6 ospf neighbor

Total number of full neighbors: 3
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.12       1   Full/Backup     00:00:32    eth1       0
10.10.10.10       1   Full/Backup     00:00:36    eth2       0
10.10.10.13      10   Full/DR         00:00:32    eth2       0
R2#

R2#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.11) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.3         10.10.10.11      945       0x80000001 0x9d7f      1
0.0.0.3         10.10.10.12      797       0x80000001 0x271c      1

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10     1766       0x80000002 0x9a8b      1
0.0.0.4         10.10.10.11     1719       0x80000002 0x3fb0      1
0.0.0.4         10.10.10.13        6       0x80000004 0xd7e9      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10     1703       0x80000005 0xf0ad      1
0.0.0.0         10.10.10.11      945       0x80000006 0xebaf      1
```

```
0.0.0.0         10.10.10.13     1708       0x80000005 0xdebc      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.4         10.10.10.13     1708        0x80000003 0xa4b1

                Inter-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.1         10.10.10.11      945        0x80000001 0xdc9f

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.13     1708        0x80000003 0xd741      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.4         10.10.10.10     1703        0x80000004 0x4ef9
0.0.0.4         10.10.10.11     1704        0x80000004 0x7acb
0.0.0.4         10.10.10.13     1708        0x80000004 0xd26f

                Router-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age        Seq#       CkSum    Link
0.0.0.0         10.10.10.11      785        0x80000003 0xa5fc      1
0.0.0.0         10.10.10.12      785        0x80000003 0x9c06      1

                Network-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.3         10.10.10.11      785        0x80000001 0x1672

                Inter-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.1         10.10.10.11      940        0x80000002 0xccaf

                Intra-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age        Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.11      784        0x80000001 0xd747      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age        Seq#       CkSum
0.0.0.3         10.10.10.11      785        0x80000002 0x70da
0.0.0.3         10.10.10.12      785        0x80000002 0x0146
```

```
R2#

R2#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10            1          10.10.10.10          eth2
10.10.10.11         B  --
10.10.10.13            1          10.10.10.13          eth2

OSPFv3 paths to Area (0.0.0.1) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.11         B  --
10.10.10.12            1          10.10.10.12          eth1

R2#

R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:09:54
C      3ffe:10::/64 via ::, eth2, 01:06:27
C      3ffe:11::/64 via ::, eth1, 00:21:16 R2#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                  Metric
     Next-hop
C  3ffe:10::/64                                     1
     directly connected, eth2, Area 0.0.0.0
C  3ffe:11::/64                                     1
     directly connected, eth1, Area 0.0.0.1
R2#


C      fe80::/64 via ::, eth9, 01:09:54
R2#
```

R1

```
R1#show ipv6 route
```

```
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:27:52
C      3ffe:10::/64 via ::, eth2, 01:25:13
O IA   3ffe:11::/64 [110/2] via fe80::5054:ff:fe3d:e317, eth2, 00:36:07
C      fe80::/64 via ::, eth9, 01:27:52
R1#
```

R3

```
R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:26:53
C      3ffe:10::/64 via ::, eth2, 01:23:21
O IA   3ffe:11::/64 [110/2] via fe80::5054:ff:fe3d:e317, eth2, 00:34:39
C      fe80::/64 via ::, eth9, 01:26:53
R3#
```

R4

```
R4#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 00:47:25
O IA   3ffe:10::/64 [110/2] via fe80::5054:ff:fe0e:46b7, eth1, 00:30:12
C      3ffe:11::/64 via ::, eth1, 00:36:23
C      fe80::/64 via ::, eth9, 00:47:25
R4#
```

© 2025 IP Infusion Inc. Proprietary

Redistribute Routes into OSPFv3

In this example, the BGP routes are imported into the OSPF routing table, and advertised as Type 5 External LSAs into Area 0.

Topology

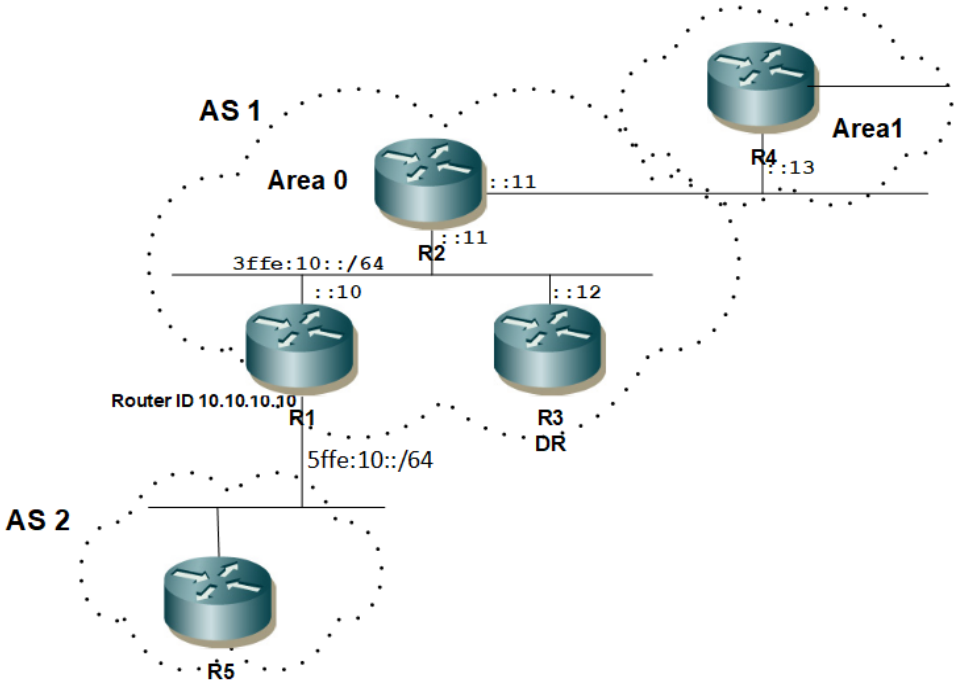

Figure 3-49: OSPFv3 Redistribute Routes

Configuration

R5

(config)#router bgp 2	Configure router bgp instance
(config-router)# neighbor 5ffe:10::10 remote-as 1	Configure R1 as ipv6 BGP neighbor
(config-router)# neighbor 5ffe:10::10 ebgp-multihop 4	Configure the ebgp -multihop for the ebgp neighbor R1
(config-router)# address-family ipv6 unicast	
(config-router-af)# neighbor 5ffe:10::10 activate	Activate the BGP neighbor on R1 for address-family ipv6 unicast to advertise and receive ipv6 routes
(config-router-af)#exit	Exit mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

R1

#configure terminal	Enter configure mode.
(config)#router bgp 2	Configure router bgp instance
(config-router)# neighbor 5ffe:10::55 remote-as 2	Configure R5 as ipv6 BGP neighbor
(config-router)# neighbor 5ffe:10::55 ebgp-multihop 4	Configure the ebgp -multihop for the ebgp neighbor R5
(config-router)# address-family ipv6 unicast	
(config-router-af)# neighbor 5ffe:10::55 activate	Activate the BGP neighbor on R5 for address-family ipv6 unicast to advertise and receive ipv6 routes
(config-router-af)#exit	Exit address-family ipv6 unicast mode
(config-router)#exit	Exit router bgp mode
(config)#router ipv6 ospf	Create an OSPFv3 routing instance.
(config-router)#router-id 10.10.10.10	Specify a Router ID (10.10.10.10) for the OSPFv3 routing process.
(config-router)#redistribute bgp	Specify redistributing routes from the other routing protocol (BGP) into OSPFv3.
(config-router)#exit	Exit OSPF router mode.
(config)#interface eth12	Enter interface mode.
(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on an interface, and assign the Area ID (0).
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R2

```
R2#show ipv6 ospf neighbor

Total number of full neighbors: 3
```

```
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.12       1   Full/Backup     00:00:31    eth1       0
10.10.10.10       1   Full/Backup     00:00:32    eth2       0
10.10.10.13      10   Full/DR         00:00:31    eth2       0
R2#

R2#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.11) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.3         10.10.10.11     1327       0x80000006 0x9384      1
0.0.0.3         10.10.10.12     1180       0x80000006 0x1d21      1

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.10.10      348       0x80000008 0x8e91      1
0.0.0.4         10.10.10.11      300       0x80000008 0x33b6      1
0.0.0.4         10.10.10.13      387       0x80000009 0xcdee      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10       24       0x8000000d 0xe6ad      1
0.0.0.0         10.10.10.11     1321       0x8000000b 0xe1b4      1
0.0.0.0         10.10.10.13      287       0x8000000b 0xd2c2      1

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.13      287       0x80000009 0x98b7

                Inter-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.1         10.10.10.11     1321       0x80000006 0xd2a4

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.13      287       0x80000009 0xcb47      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.4         10.10.10.10      284       0x8000000a 0x42ff
```

```
0.0.0.4         10.10.10.11      281       0x8000000a 0x6ed1
0.0.0.4         10.10.10.13      287       0x8000000a 0xc675

                Router-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.11     1165       0x80000008 0x9b02      1
0.0.0.0         10.10.10.12     1162       0x80000008 0x920b      1

                Network-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.11     1165       0x80000006 0x0c77

                Inter-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.1         10.10.10.11     1320       0x80000007 0xc2b4

                Inter-Area-Router-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.1         10.10.10.11       25       0x80000005 0x941a

                Intra-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.2         10.10.10.11     1165       0x80000006 0xcd4c      1  Network-LSA

                Intra-Area-Te-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.11     1165       0x80000007 0x66df
0.0.0.3         10.10.10.12     1162       0x80000007 0xf64b

                AS-external-LSA

Link State ID   ADV Router      Age       Seq#       CkSum   Route  Tag
0.0.0.1         10.10.10.10       65       0x80000002 0x284a E2      0

R2#

R2#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10        E   1          10.10.10.10          eth2
10.10.10.11         B  --
10.10.10.13            1          10.10.10.13          eth2
```

```
OSPFv3 paths to Area (0.0.0.1) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.11         B  --
10.10.10.12            1          10.10.10.12          eth1

R2#

R2#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
E2 2000::/64                                        1/20
     via fe80::5054:ff:fe2b:20b7, eth2
C  3ffe:10::/64                                      1
     directly connected, eth2, Area 0.0.0.0
C  3ffe:11::/64                                      1
     directly connected, eth1, Area 0.0.0.1
R2#

R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 03:49:59
O E2   2000::/64 [110/20] via fe80::5054:ff:fe2b:20b7, eth2, 00:36:38
C      3ffe:10::/64 via ::, eth2, 03:46:32
C      3ffe:11::/64 via ::, eth1, 03:01:21
C      fe80::/64 via ::, eth9, 03:49:59
R2#
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
```

```
     Next-hop
E2 2000::/64                                          1/20
     via fe80::5054:ff:fe2b:20b7, eth2
C  3ffe:10::/64                                         1
     directly connected, eth2, Area 0.0.0.0
IA 3ffe:11::/64                                         2
     via fe80::5054:ff:fe3d:e317, eth2, Area 0.0.0.0
R3#

R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 03:51:07
O E2   2000::/64 [110/20] via fe80::5054:ff:fe2b:20b7, eth2, 00:37:50
C      3ffe:10::/64 via ::, eth2, 03:47:35
O IA   3ffe:11::/64 [110/2] via fe80::5054:ff:fe3d:e317, eth2, 02:58:53
C      fe80::/64 via ::, eth9, 03:51:07
R3#
```

R4

R4#show ipv6 ospf route

OSPFv3 Process (*null*)

Codes: C - connected, D - Discard, 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

Destination Metric

Next-hop

E2 2000::/64 2/20

via fe80::5054:ff:fe0e:46b7, eth1

IA 3ffe:10::/64 2

via fe80::5054:ff:fe0e:46b7, eth1, Area 0.0.0.1

C 3ffe:11::/64 1

directly connected, eth1, Area 0.0.0.1

R4#

```
R4#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
```

```
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 03:15:06
O E2    2000::/64 [110/20] via fe80::5054:ff:fe0e:46b7, eth1, 00:39:34
O IA    3ffe:10::/64 [110/2] via fe80::5054:ff:fe0e:46b7, eth1, 02:57:53
C       3ffe:11::/64 via ::, eth1, 03:04:04
C       fe80::/64 via ::, eth9, 03:15:06
R4#
```

Cost

Make a route the preferred route by changing its cost. In this example, cost has been configured to make R2 the next hop for R1.

The default cost for each interface is 10. Interface eth2 on R2 has a cost of 100, and Interface eth2 on R3 has a cost of 150. The total cost to reach 10.10.14.0/24 (R4) through R2 and R3 is computed as follows:

R2: 10+100 = 110

R3: 10+150 = 160

For this reason, R1 chooses R2 as its next hop to destination 10.10.14.0/24, because it has the lower cost.

Topology

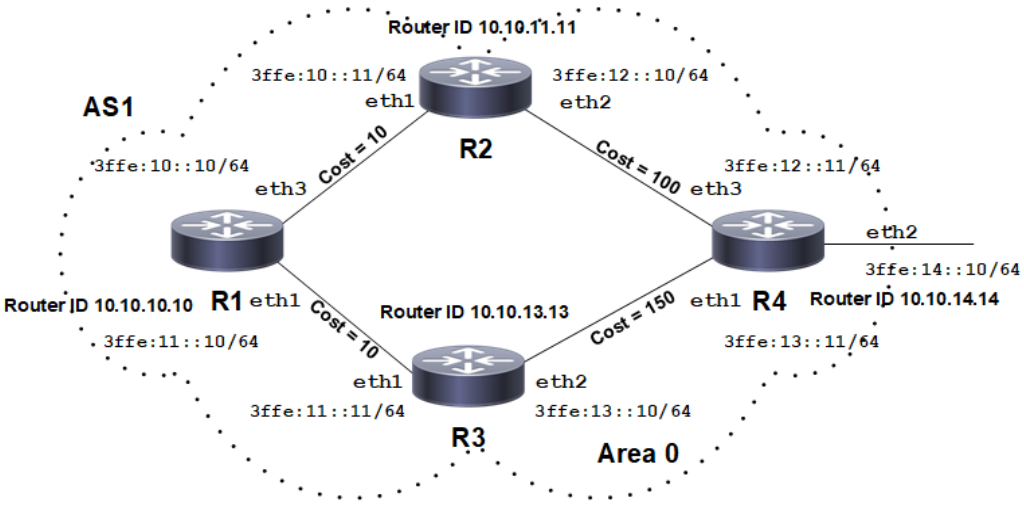

Figure 3-50: Configure Cost OSPFv3

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.10.10`	Specify a Router ID (`10.10.10.10`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth3`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.11.11`	Specify a Router ID (`10.10.11.11`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#ipv6 ospf cost 100`	Set the cost of the link-state metric (on `eth2`) to `100`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R3

`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.13.13`	Specify a Router ID (`10.10.13.13`) for the OSPFv3 routing process.

`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#ipv6 ospf cost 150`	Set the cost of link-state metric to `150`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R4

`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 10.10.14.14`	Specify a Router ID (`10.10.14.14`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth3`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID `(0)`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.13.13       1   Full/Backup     00:00:37    eth1       0
10.10.11.11       1   Full/Backup     00:00:34    eth3       0


R1#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10           --
10.10.11.11           1          10.10.11.11          eth3
10.10.13.13           1          10.10.13.13          eth1
10.10.14.14           101        10.10.11.11          eth3
```

```
rtr1#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.10) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.3         10.10.10.10      868       0x80000003 0x4839      1
0.0.0.3         10.10.13.13      747       0x80000003 0x5544      1

                Link-LSA (Interface eth3)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.5         10.10.10.10      898       0x80000003 0xf33e      1
0.0.0.3         10.10.11.11      817       0x80000003 0xce7b      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10       58       0x80000008 0xabaf      2
0.0.0.0         10.10.11.11     1767       0x80000008 0x26cd      2
0.0.0.0         10.10.13.13     1753       0x80000008 0x9724      2
0.0.0.0         10.10.14.14     1753       0x80000007 0x96b5      2

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10       58       0x80000003 0x4341
0.0.0.5         10.10.10.10      163       0x80000003 0xf88d
0.0.0.4         10.10.11.11     1767       0x80000002 0x5c22
0.0.0.4         10.10.13.13     1753       0x80000002 0x680e

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.1         10.10.10.10      813       0x80000003 0xd34b      1  Network-
LSA
0.0.0.2         10.10.10.10      743       0x80000003 0xcb53      1  Network-
LSA
0.0.0.2         10.10.11.11      652       0x80000003 0xf91f      1  Network-
LSA
0.0.0.3         10.10.13.13      684       0x80000003 0x22ec      1  Network-
LSA
                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10      868       0x80000004 0x4fe8
0.0.0.5         10.10.10.10      898       0x80000004 0x39fb
0.0.0.3         10.10.11.11      817       0x80000004 0x72c1
0.0.0.4         10.10.11.11      802       0x80000005 0xe1ea
0.0.0.3         10.10.13.13      747       0x80000004 0x5ad6
0.0.0.4         10.10.13.13      727       0x80000005 0x8f02
0.0.0.3         10.10.14.14      688       0x80000004 0x2df8
0.0.0.5         10.10.14.14      653       0x80000004 0x9c8c
```

```
rtr1#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
C  3ffe:10::/64                                        1
     directly connected, eth3, Area 0.0.0.0
C  3ffe:11::/64                                        1
     directly connected, eth1, Area 0.0.0.0
O  3ffe:12::/64                                      101
     via fe80::a00:27ff:fef9:2432, eth3, Area 0.0.0.0
O  3ffe:13::/64                                      102
     via fe80::a00:27ff:fef9:2432, eth3, Area 0.0.0.0


rtr1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:22:59
C      3ffe:10::/64 via ::, eth3, 00:51:14
C      3ffe:11::/64 via ::, eth1, 00:50:44
O      3ffe:12::/64 [110/101] via fe80::a00:27ff:fef9:2432, eth3, 00:49:33
O      3ffe:13::/64 [110/102] via fe80::a00:27ff:fef9:2432, eth3, 00:48:21
C      fe80::/64 via ::, eth1, 01:13:13
K      ff00::/8 [0/256] via ::, eth0, 01:22:47
```

R2

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/DR         00:00:32    eth1       0
10.10.14.14       1   Full/Backup     00:00:33    eth2       0


R2#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop            Interface
10.10.10.10            1          10.10.10.10         eth1
10.10.11.11            --
```

```
10.10.13.13           2          10.10.10.10          eth1
10.10.14.14           100        10.10.14.14          eth2


R2#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.11.11) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.5         10.10.10.10     1373       0x80000003 0xf33e      1
0.0.0.3         10.10.11.11     1290       0x80000003 0xce7b      1

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.11.11     1275       0x80000003 0x802a      1
0.0.0.5         10.10.14.14     1126       0x80000003 0x4f29      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10      533       0x80000008 0xabaf      2
0.0.0.0         10.10.11.11      440       0x80000009 0x24ce      2
0.0.0.0         10.10.13.13      427       0x80000009 0x9525      2
0.0.0.0         10.10.14.14      426       0x80000008 0x94b6      2

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10      533       0x80000003 0x4341
0.0.0.5         10.10.10.10      638       0x80000003 0xf88d
0.0.0.4         10.10.11.11      440       0x80000003 0x5a23
0.0.0.4         10.10.13.13      427       0x80000003 0x660f

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.1         10.10.10.10     1288       0x80000003 0xd34b      1  Network-
LSA
0.0.0.2         10.10.10.10     1218       0x80000003 0xcb53      1  Network-
LSA
0.0.0.2         10.10.11.11     1125       0x80000003 0xf91f      1  Network-
LSA
0.0.0.3         10.10.13.13     1158       0x80000003 0x22ec      1  Network-
LSA
                Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10     1343       0x80000004 0x4fe8
0.0.0.5         10.10.10.10     1373       0x80000004 0x39fb
0.0.0.3         10.10.11.11     1290       0x80000004 0x72c1
0.0.0.4         10.10.11.11     1275       0x80000005 0xe1ea
0.0.0.3         10.10.13.13     1223       0x80000004 0x5ad6
0.0.0.4         10.10.13.13     1203       0x80000005 0x8f02
```

```
0.0.0.3         10.10.14.14     1161       0x80000004 0x2df8
0.0.0.5         10.10.14.14     1126       0x80000004 0x9c8c


R2#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
C  3ffe:10::/64                                        1
     directly connected, eth1, Area 0.0.0.0
O  3ffe:11::/64                                        2
     via fe80::a00:27ff:fe6e:21d8, eth1, Area 0.0.0.0
C  3ffe:12::/64                                      100
     directly connected, eth2, Area 0.0.0.0
O  3ffe:13::/64                                      101
     via fe80::a00:27ff:fe01:c94d, eth2, Area 0.0.0.0


R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:26:25
C      3ffe:10::/64 via ::, eth1, 00:54:14
O      3ffe:11::/64 [110/2] via fe80::a00:27ff:fe6e:21d8, eth1, 00:55:03
C      3ffe:12::/64 via ::, eth2, 00:53:58
O      3ffe:13::/64 [110/101] via fe80::a00:27ff:fe01:c94d, eth2, 00:52:43
C      fe80::/64 via ::, eth2, 01:20:38
K      ff00::/8 [0/256] via ::, eth2, 01:20:39
```

R3

```
R3#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/DR         00:00:33    eth1       0
10.10.14.14       1   Full/Backup     00:00:38    eth2       0



R3#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop             Interface
10.10.10.10            1          10.10.10.10          eth1
10.10.11.11            2          10.10.10.10          eth1
10.10.13.13            --
```

```
10.10.14.14            102        10.10.10.10          eth1


R3#
R3#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.13.13) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.3         10.10.10.10     1591       0x80000003 0x4839      1
0.0.0.3         10.10.13.13     1468       0x80000003 0x5544      1

                Link-LSA (Interface eth2)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.13.13     1448       0x80000003 0x9d29      1
0.0.0.3         10.10.14.14     1409       0x80000003 0x50cf      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10      780       0x80000008 0xabaf      2
0.0.0.0         10.10.11.11      689       0x80000009 0x24ce      2
0.0.0.0         10.10.13.13      673       0x80000009 0x9525      2
0.0.0.0         10.10.14.14      673       0x80000008 0x94b6      2

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10      780       0x80000003 0x4341
0.0.0.5         10.10.10.10      885       0x80000003 0xf88d
0.0.0.4         10.10.11.11      689       0x80000003 0x5a23
0.0.0.4         10.10.13.13      673       0x80000003 0x660f

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.1         10.10.10.10     1536       0x80000003 0xd34b      1  Network-
LSA
0.0.0.2         10.10.10.10     1466       0x80000003 0xcb53      1  Network-
LSA
0.0.0.2         10.10.11.11     1374       0x80000003 0xf91f      1  Network-
LSA
0.0.0.3         10.10.13.13     1403       0x80000003 0x22ec      1  Network-
LSA
               Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10     1591       0x80000004 0x4fe8
0.0.0.5         10.10.10.10     1621       0x80000004 0x39fb
0.0.0.3         10.10.11.11     1539       0x80000004 0x72c1
0.0.0.4         10.10.11.11     1524       0x80000005 0xe1ea
0.0.0.3         10.10.13.13     1468       0x80000004 0x5ad6
0.0.0.4         10.10.13.13     1448       0x80000005 0x8f02
0.0.0.3         10.10.14.14     1409       0x80000004 0x2df8
```

```
0.0.0.5         10.10.14.14     1374       0x80000004 0x9c8c


R3#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
O  3ffe:10::/64                                        2
     via fe80::a00:27ff:fe7d:2a72, eth1, Area 0.0.0.0
C  3ffe:11::/64                                        1
     directly connected, eth1, Area 0.0.0.0
O  3ffe:12::/64                                      102
     via fe80::a00:27ff:fe7d:2a72, eth1, Area 0.0.0.0
O  3ffe:13::/64                                      103
     via fe80::a00:27ff:fe7d:2a72, eth1, Area 0.0.0.0


R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:28:16
O      3ffe:10::/64 [110/2] via fe80::a00:27ff:fe7d:2a72, eth1, 00:58:14
C      3ffe:11::/64 via ::, eth1, 00:55:44
O      3ffe:12::/64 [110/102] via fe80::a00:27ff:fe7d:2a72, eth1, 00:56:36
C      3ffe:13::/64 via ::, eth2, 00:55:26
C      fe80::/64 via ::, eth1, 01:20:39
K      ff00::/8 [0/256] via ::, eth2, 01:21:40
```

R4

```
R4#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.13.13       1   Full/DR         00:00:30    eth1       0
10.10.11.11       1   Full/DR         00:00:30    eth3       0



R4#show ipv6 ospfv3 topology

OSPFv3 Process (*null*)
OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop            Interface
10.10.10.10             2          10.10.11.11         eth3
                                   10.10.13.13         eth1
10.10.11.11             1          10.10.11.11         eth3
```

```
10.10.13.13            1          10.10.13.13          eth1
10.10.14.14            --


R4#show ipv6 ospf database

            OSPFv3 Router with ID (10.10.14.14) (Process *null*)

                Link-LSA (Interface eth1)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.13.13     1634       0x80000003 0x9d29      1
0.0.0.3         10.10.14.14     1592       0x80000003 0x50cf      1

                Link-LSA (Interface eth3)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix
0.0.0.4         10.10.11.11     1708       0x80000003 0x802a      1
0.0.0.5         10.10.14.14     1557       0x80000003 0x4f29      1

                Router-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum    Link
0.0.0.0         10.10.10.10      966       0x80000008 0xabaf      2
0.0.0.0         10.10.11.11      873       0x80000009 0x24ce      2
0.0.0.0         10.10.13.13      859       0x80000009 0x9525      2
0.0.0.0         10.10.14.14      857       0x80000008 0x94b6      2

                Network-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10      966       0x80000003 0x4341
0.0.0.5         10.10.10.10     1071       0x80000003 0xf88d
0.0.0.4         10.10.11.11      873       0x80000003 0x5a23
0.0.0.4         10.10.13.13      859       0x80000003 0x660f

                Intra-Area-Prefix-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum  Prefix  Reference
0.0.0.1         10.10.10.10     1721       0x80000003 0xd34b      1  Network-
LSA
0.0.0.2         10.10.10.10     1651       0x80000003 0xcb53      1  Network-
LSA
0.0.0.2         10.10.11.11     1558       0x80000003 0xf91f      1  Network-
LSA
0.0.0.3         10.10.13.13     1589       0x80000003 0x22ec      1  Network-
LSA

  Intra-Area-Te-LSA (Area 0.0.0.0)

Link State ID   ADV Router      Age       Seq#       CkSum
0.0.0.3         10.10.10.10     1776       0x80000004 0x4fe8
0.0.0.5         10.10.10.10        6       0x80000005 0x37fc
0.0.0.3         10.10.11.11     1723       0x80000004 0x72c1
0.0.0.4         10.10.11.11     1708       0x80000005 0xe1ea
0.0.0.3         10.10.13.13     1654       0x80000004 0x5ad6
0.0.0.4         10.10.13.13     1634       0x80000005 0x8f02
```

```
0.0.0.3         10.10.14.14     1592       0x80000004 0x2df8
0.0.0.5         10.10.14.14     1557       0x80000004 0x9c8c


R4#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
O  3ffe:10::/64                                        2
     via fe80::a00:27ff:fe0d:fbe3, eth3, Area 0.0.0.0
O  3ffe:11::/64                                        2
     via fe80::a00:27ff:fecf:8873, eth1, Area 0.0.0.0
C  3ffe:12::/64                                        1
     directly connected, eth3, Area 0.0.0.0
C  3ffe:13::/64                                        1
     directly connected, eth1, Area 0.0.0.0


R4#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:32:01
O      3ffe:10::/64 [110/2] via fe80::a00:27ff:fe0d:fbe3, eth3, 01:02:49
O      3ffe:11::/64 [110/2] via fe80::a00:27ff:fecf:8873, eth1, 01:02:19
C      3ffe:12::/64 via ::, eth3, 00:58:46
C      3ffe:13::/64 via ::, eth1, 00:59:18
C      fe80::/64 via ::, eth1, 01:27:01
K      ff00::/8 [0/256] via ::, eth3, 01:27:31
```

Virtual Links

Virtual links are used to connect a temporarily-disjointed non-backbone area to the backbone area, or to repair a non-contiguous backbone area. In this example, the ABR R3 has temporarily lost connection to Area 0, in turn disconnecting Area 2 from the backbone area. The virtual link between ABR R1 and ABR R2 connects Area 2 to Area 0. Area 1 is used as a transit area.

Topology

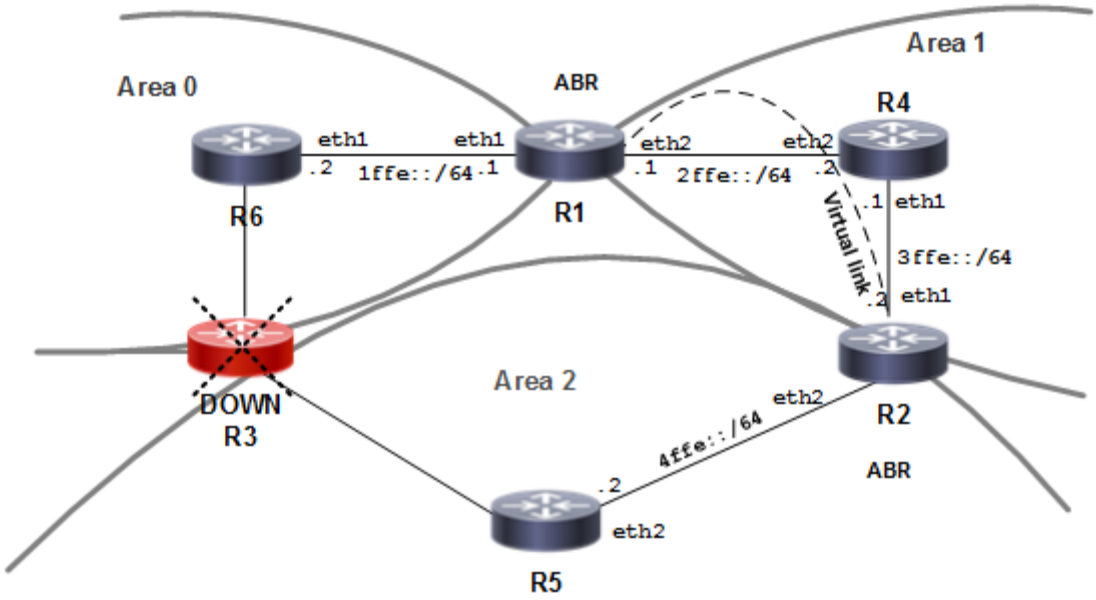

Figure 3-51: OSPFv3 Virtual Links

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Setup loopback interface
`(config-if)#ip address 1.1.1.1/32 secondary`	Specify loopback interface address
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on this interface, and assign the Area ID `(0)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on this interface, and assign the Area ID `(1)`.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 1.1.1.1`	Specify a Router ID (`1.1.1.1`) for the OSPFv3 routing process.
`(config-router)#area 1 virtual-link 2.2.2.2`	Configure a virtual link between this router R1 and R2 (Router ID `2.2.2.2`) through transit area 1.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Setup loopback interface
`(config-if)#ip address 2.2.2.2/32 secondary`	Specify loopback interface address
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on this interface, and assign the Area ID `(1)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 2`	Enable OSPFv3 routing on this interface, and assign the Area ID `(2)`.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 2.2.2.2`	Specify a Router ID (`2.2.2.2)` for the OSPFv3 routing process.
`(config-router)#area 1 virtual-link 1.1.1.1`	Configure a virtual link between this router R1 and R2 (Router ID `1.1.1.1`) through transit area 1.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R4

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on this interface, and assign the Area ID `(1)`.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 1`	Enable OSPFv3 routing on this interface, and assign the Area ID `(1)`.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf`	Create an OSPFv3 routing instance.
`(config-router)#router-id 4.4.4.4`	Specify a Router ID (`4.4.4.4`) for the OSPFv3 routing process.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R5

`#configure terminal`	Enter configure mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router ospf area 2`	Enable OSPFv3 routing on this interface, and assign the Area ID `(2)`.
`(config-if)#exit`	Exit interface mode.

(config)#router ipv6 ospf	Create an OSPFv3 routing instance.
(config-router)#router-id 5.5.5.5	Specify a Router ID (5.5.5.5) for the OSPFv3 routing process.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R6

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on this interface, and assign the Area ID (0).
(config-if)#exit	Exit interface mode.
(config)#router ipv6 ospf	Create an OSPFv3 routing instance.
(config-router)#router-id 6.6.6.6	Specify a Router ID (6.6.6.6) for the OSPFv3 routing process.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R2

```
#show ipv6 ospf n
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/DR         00:00:31    eth1       0
3.3.3.3           1   Full/DR         00:00:32    eth2       0
2.2.2.2           1   Full/ -         inactive    VLINK2147479553 0


#show ipv6 ospf virtual-links
Virtual Link VLINK2147479553 to router 2.2.2.2 is up
  Transit area 0.0.0.1 via interface eth2, instance ID 0
  Hello suppression Enabled
  DoNotAge LSA allowed
  Local address 2ffe::11/128
  Remote address 3ffe::11/128
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in inactive
    Adajcency state Full


# show ipv6 ospf
 Routing Process "OSPFv3 (*null*)" with ID 1.1.1.1
 Process uptime is 5 minutes
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 This router is an ASBR (injecting external routing information)
```

```
 SPF schedule delay initial 0.500 secs
 SPF schedule delay min 0.500 secs
 SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incoming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 0. Checksum Sum 0x0000
 Number of AS-Scoped Unknown LSA 0
 Number of LSA originated 17
 Number of LSA received 50
 Number of areas in this router is 2
    Area BACKBONE(0)
        Number of interfaces in this area is 2(2)
        SPF algorithm executed 8 times
        Number of LSA 23.  Checksum Sum 0xB35D8
        Number of Unknown LSA 0
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 13 times
        Number of LSA 16.  Checksum Sum 0x7845A
        Number of Unknown LSA 0
 Dste Staus: Disabled


#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  1ffe::/64                                         1
     directly connected, eth1, Area 0.0.0.0
C  2ffe::/64                                         1
     directly connected, eth2, TransitArea 0.0.0.1
C  2ffe::11/128                                      0
     directly connected, eth2, TransitArea 0.0.0.1
O  3ffe::/64                                         1
     directly connected, eth2, TransitArea 0.0.0.1
O  3ffe::11/128                                      2
     via fe80::5054:ff:fe6f:334d, eth2, TransitArea 0.0.0.1
IA 4ffe::/64                                         3
     via fe80::5054:ff:fe6f:334d, eth2, TransitArea 0.0.0.1
#
```

R3

```
#show ipv6 ospf n
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
```

```
1.1.1.1           1   Full/Backup     00:00:35    eth1       0
2.2.2.2           1   Full/Backup     00:00:30    eth2       0


# show ipv6 ospf
 Routing Process "OSPFv3 (*null*)" with ID 3.3.3.3
 Process uptime is 5 minutes
 This router is an ASBR (injecting external routing information)
 SPF schedule delay initial 0.500 secs
 SPF schedule delay min 0.500 secs
 SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incoming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 0. Checksum Sum 0x0000
 Number of AS-Scoped Unknown LSA 0
 Number of LSA originated 10
 Number of LSA received 23
 Number of areas in this router is 1
    Area 0.0.0.1
        Number of interfaces in this area is 2(2)
        SPF algorithm executed 14 times
        Number of LSA 16.  Checksum Sum 0x7845A
        Number of Unknown LSA 0
 Dste Staus: Disabled


#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
IA 1ffe::/64                                          2
     via fe80::5054:ff:feb7:cc69, eth1, TransitArea 0.0.0.1
C  2ffe::/64                                          1
     directly connected, eth2, TransitArea 0.0.0.1
O  2ffe::11/128                                       1
     via fe80::5054:ff:feb7:cc69, eth1, TransitArea 0.0.0.1
C  3ffe::/64                                          1
     directly connected, eth1, TransitArea 0.0.0.1
O  3ffe::11/128                                       1
     via fe80::5054:ff:fec5:2430, eth2, TransitArea 0.0.0.1
IA 4ffe::/64                                          2
     via fe80::5054:ff:fec5:2430, eth2, TransitArea 0.0.0.1
```

R4

```
#show ipv6 ospf n
```

```
OSPFv3 Process (*null*)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
3.3.3.3           1   Full/DR         00:00:31    eth1       0
1.1.1.1           1   Full/ -         inactive    VLINK2147479554 0


#show ipv6 ospf virtual-links
Virtual Link VLINK2147479554 to router 1.1.1.1 is up
  Transit area 0.0.0.1 via interface eth1, instance ID 0
  Hello suppression Enabled
  DoNotAge LSA allowed
  Local address 3ffe::11/128
  Remote address 2ffe::11/128
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in inactive
    Adajcency state Full


# show ipv6 ospf
 Routing Process "OSPFv3 (*null*)" with ID 2.2.2.2
 Process uptime is 4 minutes
 This router is an ABR, ABR Type is Alternative Cisco (RFC3509)
 This router is an ASBR (injecting external routing information)
 SPF schedule delay initial 0.500 secs
 SPF schedule delay min 0.500 secs
 SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incoming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 0. Checksum Sum 0x0000
 Number of AS-Scoped Unknown LSA 0
 Number of LSA originated 26
 Number of LSA received 37
 Number of areas in this router is 3
    Area BACKBONE(0)
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 3 times
        Number of LSA 23.  Checksum Sum 0xB35D8
        Number of Unknown LSA 0
    Area 0.0.0.1
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 11 times
        Number of LSA 16.  Checksum Sum 0x7845A
        Number of Unknown LSA 0
    Area 0.0.0.2
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 4 times
        Number of LSA 11.  Checksum Sum 0x5D8B7
        Number of Unknown LSA 0
```

```
 Dste Staus: Disabled


#show ipv6 ospf route
OSPFv3 Process (*null*)
Codes: C - connected, D - Discard, 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

   Destination                                      Metric
     Next-hop
O  1ffe::/64                                          3
     via fe80::5054:ff:fe1d:eace, eth1, TransitArea 0.0.0.1
O  2ffe::/64                                          1
     directly connected, eth1, TransitArea 0.0.0.1
O  2ffe::11/128                                       2
     via fe80::5054:ff:fe1d:eace, eth1, TransitArea 0.0.0.1
C  3ffe::/64                                          1
     directly connected, eth1, TransitArea 0.0.0.1
C  3ffe::11/128                                       0
     directly connected, eth1, TransitArea 0.0.0.1
C  4ffe::/64                                          1
     directly connected, eth2, Area 0.0.0.2
```

Multiple Instances

By using multiple OSPFv3 instances, OSPFv3 routes can be segregated, based on their instance number. Routes of one instance are stored differently from routes of another instance running in the same router.

To configure multiple OSPFv3 instances, refer to the topology diagram and follow the procedures below.

1. Enable OSPFv3 on an interface.
2. Enable multiple instances.
3. Configure redistribution among multiple instances.

   Note: Optionally, redistribution can be configured with the metric, type, or route-map options.

Topology

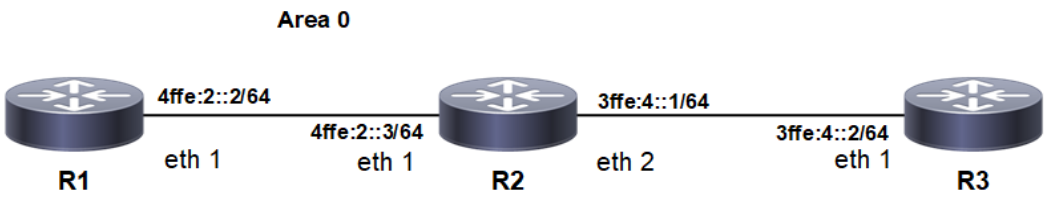

Enable Multiple OSPFv3 Instances on a Router Based on Tags

In this example, routers R1, R2, and R3 are in Area 0, and all run OSPFv3.

R1

`(config)#router ipv6 ospf 5`	Configure an OSPFv3 instance with an instance ID of `5`.
`(config-router)#router-id 5.5.5.5`	Configure the router ID to use on this instance.
`(config-router)#exit`	Exit Router mode, and return to Interface mode.
`(config)#interface eth1`	Specify the interface on which OSPFv3 is to be enabled.
`(config-if)#ipv6 address 4ffe:2::2/64`	Configure the IPv6 address.
`(config-if)#ipv6 router ospf area 0 tag 5`	Configure the area number and instance value: match the instance ID with the instance ID created previously.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R2

`(config)#router ipv6 ospf 5`	Configure an OSPFv3 instance with an instance ID of `5`.
`(config-router)#router-id 149.149.149.149`	Configure the router ID to use on this instance.
`(config-router)#exit`	Exit Router mode, and return to Interface mode.
`(config)#interface eth1`	Configure the interface to connect to R1.
`(config-if)#ipv6 address 4ffe:2::3/64`	Configure the IPv6 address.
`(config-if)#ipv6 router ospf area 0 tag 5`	Configure the area number and instance value: match the instance ID with the instance ID created previously.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit Interface mode, and return to Configure terminal mode.
`(config)#router ipv6 ospf 15`	Configure an OSPFv3 instance with an instance ID of `15`.
`(config-router)#router-id 159.159.159.159`	Configure the router ID to use on this instance.
`(config-router)#exit`	Exit Router mode, and return to Interface mode.
`(config)#interface eth2`	Configure the interface to connect to R3.
`(config-if)#ipv6 address 3ffe:4::1/64`	Configure the IPv6 address.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#ipv6 router ospf area 0 tag 15`	Configure the area number and instance value: match the instance ID with the instance ID created previously.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R3

`(config)#router ipv6 ospf 15`	Configure an OSPFv3 instance with an instance ID of `15`.
`(config-router)#router-id 152.152.152.152`	Configure the router ID to use on this instance.

`(config-router)#exit`	Exit Router mode, and return to Interface mode.
`(config)#interface eth1`	Specify the interface on which OSPFv3 is to be enabled.
`(config-if)#ipv6 address 3ffe:4::2/64`	Configure the IPv6 address.
`(config-if)#ipv6 router ospf area 0 tag 15`	Configure the area number and instance value: match the instance ID with the instance ID created previously.
`(config-if)#no shutdown`	Activate the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  4ffe:2::/64                                        1
     directly connected, eth1, Area 0.0.0.0

R1#show ipv6 ospf neighbor
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
149.149.149.149   1   Full/Backup     00:00:32    eth1       0
```

R2

```
R2#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth2, Area 0.0.0.0

OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  4ffe:2::/64                                        1
     directly connected, eth1, Area 0.0.0.0

R2#sh ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
```

```
152.152.152.152   1   Full/DR         00:00:35    eth2       0
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
5.5.5.5           1   Full/DR         00:00:33    eth1       0
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth1, Area 0.0.0.0


R3#sh ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
159.159.159.159   1   Full/Backup     00:00:34    eth1       0
```

Redistribute among Multiple Instances

In this example, routes of one ospfv3 instance are redistributed to another ospfv3 instance to enable ping from R1 to R3 or vice-versa; and R2 redistributes routes from one instance to another.

R2

(config)#router ipv6 ospf 15	Configure an OSPFv3 instance with instance ID 15.
(config-router)#router-id 159.159.159.159	Configure the router ID.
(config-router)#redistribute ospf 5	Redistribute instance 5 routes.
(config-router)#exit	Exit OSPF router mode.
(config)#router ipv6 ospf 5	Configure an OSPFv3 instance with instance ID 5.
(config-router)#router-id 149.149.149.149	Configure the router ID.
(config-router)#redistribute ospf 15	Redistribute instance 15 routes.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
```

```
E2 3ffe:4::/64                                        1/20
     via fe80::a00:27ff:fef9:2432, eth1
C  4ffe:2::/64                                           1
     directly connected, eth1, Area 0.0.0.0


rtr1#show ipv6 ospf neighbor
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
149.149.149.149   1   Full/Backup     00:00:32    eth1       0
```

R2

```
R2#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                        Metric
     Next-hop
C  3ffe:4::/64                                           1
     directly connected, eth2, Area 0.0.0.0
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                        Metric
     Next-hop
C  4ffe:2::/64                                           1
     directly connected, eth1, Area 0.0.0.0

R2#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
152.152.152.152   1   Full/DR         00:00:34    eth2       0
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
5.5.5.5           1   Full/DR         00:00:30    eth1       0
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                        Metric
     Next-hop
C  3ffe:4::/64                                           1
     directly connected, eth1, Area 0.0.0.0
E2 4ffe:2::/64                                        1/20
     via fe80::a00:27ff:fe0d:fbe3, eth1
```

```
R3#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
159.159.159.159   1   Full/Backup     00:00:33    eth1       0
```

Redistribute with Metric Option

In this example, on R3, routes of instance 15 are redistributed into instance and vice-versa with metric of 100 so that R1 and R2 have each other's routes with a metric of 100.

R2

`(config)#router ipv6 ospf 15`	Configure an OSPFv3 instance with instance ID `15`.
`(config-router)#router-id 159.159.159.159`	Configure the router ID.
`(config-router)#redistribute ospf 5 metric 100`	Redistribute instance `5` routes with metric `100`.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#router ipv6 ospf 5`	Redistribute routes into instance `5`.
`(config-router)#router-id 149.149.149.149`	Configure the router ID.
`(config-router)#redistribute ospf 15 metric 100`	Redistribute instance `15` routes with metric `100`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
E2 3ffe:4::/64                                      1/100
     via fe80::a00:27ff:fef9:2432, eth1
C  4ffe:2::/64                                         1
     directly connected, eth1, Area 0.0.0.0
R1#

R1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime
```

```
IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 18:08:02
O E2    3ffe:4::/64 [110/100] via fe80::5054:ff:fe0e:46b7, eth1, 00:00:25
C       4ffe:2::/64 via ::, eth1, 00:20:39
C       fe80::/64 via ::, eth9, 18:08:02
R1#


R1#show ipv6 ospf neighbor
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
149.149.149.149   1   Full/Backup     00:00:36    eth1       0
```

R2

```
R2#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth2, Area 0.0.0.0
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  4ffe:2::/64                                        1
     directly connected, eth1, Area 0.0.0.0



R2#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
152.152.152.152   1   Full/DR         00:00:33    eth2       0
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
5.5.5.5           1   Full/DR         00:00:40    eth1       0
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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
```

```
   Destination                                        Metric
     Next-hop
C  3ffe:4::/64                                           1
     directly connected, eth1, Area 0.0.0.0
E2 4ffe:2::/64                                         1/100
     via fe80::a00:27ff:fe0d:fbe3, eth1

R3#

R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 18:08:13
C      3ffe:4::/64 via ::, eth1, 00:17:55
O E2   4ffe:2::/64 [110/100] via fe80::5054:ff:fe3d:e317, eth1, 00:01:05
C      fe80::/64 via ::, eth9, 18:08:13
R3#


R3#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
159.159.159.159   1   Full/Backup     00:00:37    eth1       0
```

Redistribute with Type Option

In this example, on R3, R1 has R3 routes as type 2, and R3 has R1 routes as type 1.

R2

`(config)#router ipv6 ospf 15`	Configure an OSPFv3 instance with instance ID `15`.
`(config-router)#router-id 159.159.159.159`	Configure the router ID.
`(config-router)#redistribute ospf 5 metric-type 1`	Redistribute instance `5` routes as type `1`.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#router ipv6 ospf 5`	Redistribute routes into instance `5`.
`(config-router)#router-id 149.149.149.149`	Configure the router ID.
`(config-router)#redistribute ospf 15 metric-type 2`	Redistribute instance `15` routes as type `2`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
E2 3ffe:4::/64                                      1/20
     via fe80::a00:27ff:fef9:2432, eth1
C  4ffe:2::/64                                         1
     directly connected, eth1, Area 0.0.0.0


rtr1#show ipv6 ospf neighbor
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance I
D
149.149.149.149   1   Full/Backup     00:00:32    eth1       0
```

R2

```
R2#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                         1
     directly connected, eth2, Area 0.0.0.0
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  4ffe:2::/64                                         1
     directly connected, eth1, Area 0.0.0.0

R2#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance I
D
152.152.152.152   1   Full/DR         00:00:36    eth2       0
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance I
D
5.5.5.5           1   Full/DR         00:00:32    eth1       0
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth1, Area 0.0.0.0
E2 4ffe:2::/64                                      1/21
     via fe80::a00:27ff:fe0d:fbe3, eth1


R3#show ipv6 ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance I
D
159.159.159.159   1   Full/Backup     00:00:36    eth1       0
```

Redistribute with Route-Map Option

R1

`(config)#interface eth2`	Configure the interface eth2 on R1.
`(config-if)#ipv6 address 4ffe:1::2/64`	Configure the IPv6 address.
`(config-if)#ipv6 router ospf area 0 tag 5`	Configure interface eth2 for ospfv3 with area 0 and instance 5
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

Validation

```
R3
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                         1
     directly connected, eth1, Area 0.0.0.0
E1 4ffe:1::/64                                        20
     via fe80::5054:ff:fe3d:e317, eth1
E1 4ffe:2::/64                                        20
     via fe80::5054:ff:fe3d:e317, eth1
R3#
```

R2

`R2(config)#ipv6 prefix-list permit-4ffe-2`	Configure and ipv6 prefix-list
`R2 (config-ipv6-prefix-list)#seq 5 permit 4ffe:2::/64`	Create an access-list to permit the prefix 4ffe:2::/64
`R2(config-ipv6-prefix-list)#exit`	Exit the prefix-list mode
`R2(config)#route-map permit-only-4ffe-2`	Configure a route-map to permit only the prefix 4ffe:2::/64
`R2(config-route-map)#match ipv6 address prefix-list permit-4ffe-2`	Configure a match statement to match the configured ipv6 prefix-list
`R2(config-route-map)#exit`	Exit route-map mode and return to configure terminal mode
`R2(config)#router ipv6 ospf 15`	Enter router ipv6 mode for instance 15
`R2(config-router)#redistribute ospf 5 route-map permit-only-4ffe-2`	Redistribute instance 5 routes with route-map to permit only the ipv6 prefix 4ffe:2::/64
`(config-router)#exit`	Exit OSPF router mode.
`(config)#router ipv6 ospf 5`	Redistribute routes into instance 5.
`(config-router)#router-id 149.149.149.149`	Configure the router ID.
`(config-router)#redistribute ospf 15 route-map 1`	Redistribute instance 15 routes with route map 1.
`(config-router)#redistribute connected`	Redistribute connected routes to instance 15.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
E2 3ffe:4::/64                                      1/20
     via fe80::5054:ff:fe0e:46b7, eth1
C  4ffe:1::/64                                         1
     directly connected, eth2, Area 0.0.0.0
C  4ffe:2::/64                                         1
     directly connected, eth1, Area 0.0.0.0
R1#

R1#show ipv6 ospf neighbor
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance I
D
149.149.149.149   1   Full/DR         00:00:34    eth1       0
```

R2

```
R2#show ipv6 ospf route
```

```
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth2, Area 0.0.0.0
OSPFv3 Process (5)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
O  4ffe:1::/64                                        2
     via fe80::5054:ff:fe0c:40ed, eth1, Area 0.0.0.0
C  4ffe:2::/64                                        1
     directly connected, eth1, Area 0.0.0.0
R2#


R2#sh ipv6  ospf neighbor
OSPFv3 Process (15)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
152.152.152.152   1   Full/Backup     00:00:32    eth2       0
OSPFv3 Process (5)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
5.5.5.5           1   Full/Backup     00:00:38    eth1       0
```

R3

```
R3#show ipv6 ospf route
OSPFv3 Process (15)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  3ffe:4::/64                                        1
     directly connected, eth1, Area 0.0.0.0
E2 4ffe:2::/64                                     1/20
     via fe80::5054:ff:fe3d:e317, eth1
R3#
```

Not-So-Stubby Area

This section contains OSPFv3 NSSA (Not-So-Stubby Area) configuration examples.

An NSSA allows external routes to be advertised into the OSPF autonomous system while retaining the characteristics of a stub area to the rest of the autonomous system. To do this, the ASBR in an NSSA will originate type 7 LSAs to

advertise the external destinations. These NSSA external LSAs are flooded throughout the NSSA but are blocked at the ABR.

The NSSA external LSA has a flag in its header known as the P-bit. The NSSA ASBR has the option of setting or clearing the P-bit. If an NSSA's ABR receives a type 7 LSA with the P-bit set to one, it translates the type 7 LSA into a type 5 LSA and floods it throughout the other areas. If the P-bit is set to zero, no translation takes place and the destination in the type 7 LSA is not advertised outside of the NSSA.

Topology

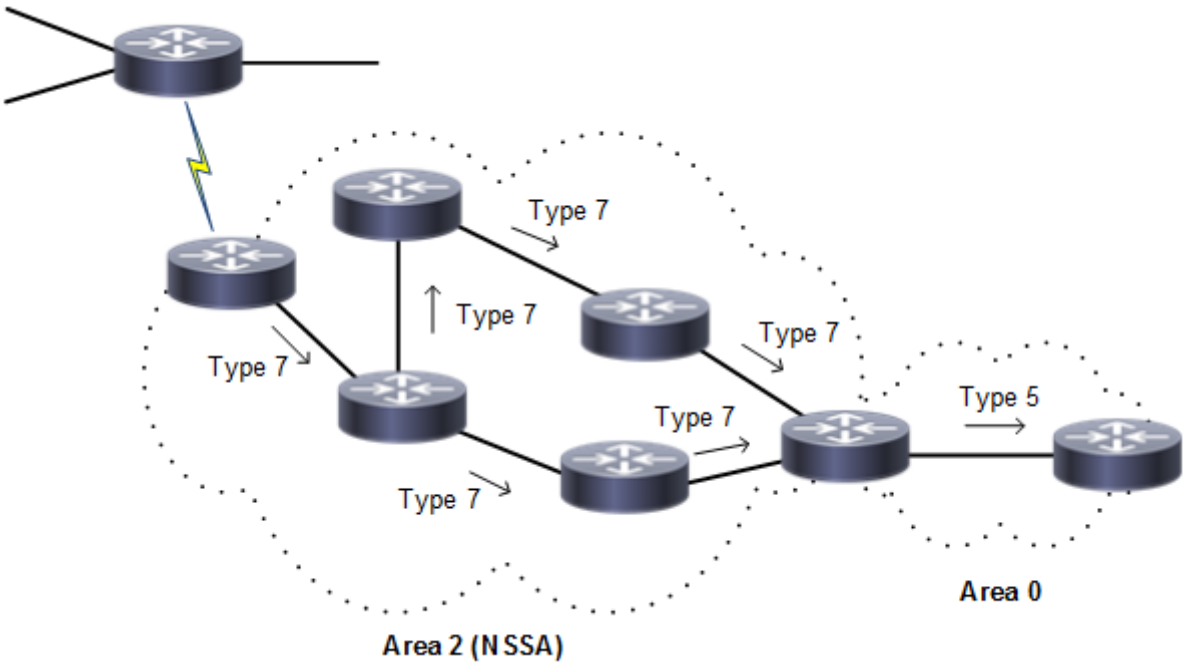

Figure 3-52: Translating Type 7 LSAs into Type 5 LSAs

NSSA with Route Option

This example shows the configuration to enable NSSA and to configure different route options for NSSA. There are three `area nssa` command options for originating default Type-3 LSA and default Type-7 LSA, and for blocking redistribution of Type-7 LSA into an NSSA:

- `no-summary`: The NSSA ABR blocks all type-3 and type-4 LSAs into the NSSA area and sends a single type-3 LSA into the area to advertise a default route
- `default-information-originate`: The NSSA ABR advertises a default route into the NSSA as a type-7 LSA.
- `no-redistribution`: The NSSA ABR bocks type-7 LSA from being redistributed into the NSSA area.

In Figure 3-53, R2 is an NSSA ABR as well as an NSSA ASBR that maps the router interfaces to two different areas and redistributes the connected routes of the loopback interface. Also, this example sets the `no-summary`, `no-redistribution`, and `default-information-originate` options on R2 to originate default Type-3 LSAs and default Type-7 LSAs into the NSSA and to block Type-7 LSAs.

Topology

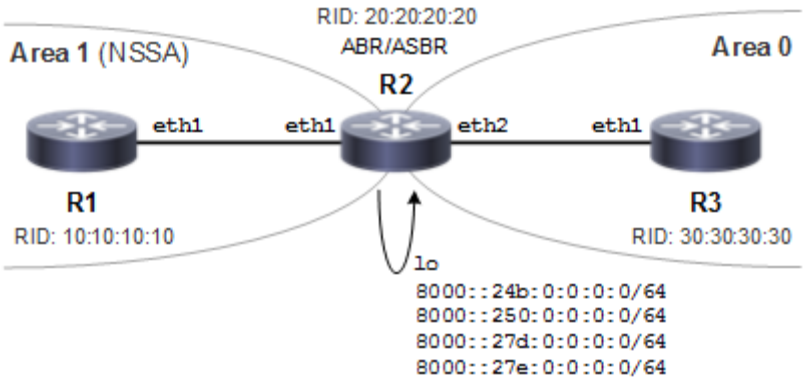

Figure 3-53: NSSA with Route Options

R1

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 10.10.10.10`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa`	Configure area as NSSA
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

R2

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode for `eth2`
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone
`(config-if)#interface lo`	Enter interface mode for Loopback
`(config-if)#ipv6 address 8000::24b:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::250:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::27d:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::27e:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.

© 2025 IP Infusion Inc. Proprietary

(config)#router ipv6 ospf 100	Configure the routing process and specify the tag (100).
(config-router)#router-id 20.20.20.20	Configure the router ID to use on this instance (100)
(config-router)#redistribute connected	Redistribute the configured loopback network into the NSSA
(config-router)#area 1 nssa no-redistribution default-information-originate no-summary	Configure the Router to originate default Type-3 LSAs and default Type-7 LSAs, and to block Type-7 LSAs into the NSSA
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit interface mode

R3

(config)#interface eth1	Enter interface mode for eth1.
(config-if)#ipv6 router ospf tag 100 area 0	Configure interface in an area assigned with the area ID (1).
(config-if)#exit	Exit interface mode.
(config)#router ipv6 ospf 100	Configure the routing process and specify the tag (100).
(config-router)#router-id 30.30.30.30	Configure the router ID to use on this instance (100)
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit interface mode

Validation 1

In the output of `show ipv6 ospf neighbor` below, verify that OSPFv3 adjacency is in state “full” for both R1 and R2 under the process identifier 100.

```
R1#sh ipv6 ospf neighbor
OSPFv3 Process (100)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
20.20.20.20       1   Full/DR         00:00:34    eth1       0

R2#show ipv6 ospf neighbor

Total number of full neighbors: 2
OSPFv3 Process (100)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
10.10.10.10       1   Full/DR         00:00:36    eth1       0
30.30.30.30       1   Full/Backup     00:00:39    eth2       0
R2#
```

Validation 2

The output below shows originating default Type-3 LSAs into the NSSA with the no-summary option. The advertising router identifier is for R2 (20.20.20.20, the NSSA-ABR). Also, the prefix is ::/0 and the LS-Type is Inter-Area-Prefix-LSA for the default Type-3 LSA route into the NSSA.

```
R1#sh ipv6 ospf database inter-prefix

            OSPFv3 Router with ID (10.10.10.10) (Process 100)
```

```
                Inter-Area-Prefix-LSA (Area 0.0.0.1)

  LS age: 1234
  LS Type: Inter-Area-Prefix-LSA
  Link State ID: 0.0.0.6
  Advertising Router: 20.20.20.20
  LS Seq Number: 0x80000001
  Checksum: 0x17D0
  Length: 28
    Metric: 1
    Prefix: ::/0
    Prefix Options: 0
```

Validation 3

The output below shows originating default type-7 LSAs alone after setting the no-redistribution and default-information originate options. The advertising router identifier is for R2 (20.20.20.20, the NSSA-ABR). Also, the prefix is ::/0 and LS-Type is NSSA-external-LSA for the default Type-7 LSA route into the NSSA

```
R1#sh ipv6 ospf database nssa-external

            OSPFv3 Router with ID (10.10.10.10) (Process 100)

                NSSA-external-LSA (Area 0.0.0.1)

  LS age: 1758
  LS Type: NSSA-external-LSA
  Link State ID: 0.0.0.20
  Advertising Router: 20.20.20.20
  LS Seq Number: 0x80000002
  Checksum: 0x6468
  Length: 32
    Metric Type: 2 (Larger than any link state path)
    Metric: 1
    Prefix: ::/0
    Prefix Options: 0 (-|-|-|-)
    External Route Tag: 0
```

NSSA with the Summary Address Option

Figure 3-54 shows the configuration to originate external LSAs (Type-7) and translate them into external LSAs (Type-5):

- R1 is an NSSA-ASBR configured with loopback IPv6 addresses that are redistributed into OSPFv3
- R2 is an NSSA-ABR
- R3 is backbone router

R1 originates Type-7 LSAs which are summarized into a single Type-7 into the NSSA by the `summary-address` option and this summarized Type-7 is converted to Type-5 LSA by R2.

Also, the summarized route can be tagged using the `tag` command and the advertisement of summarized routes can be suppressed by the `not-advertise` option.

Topology

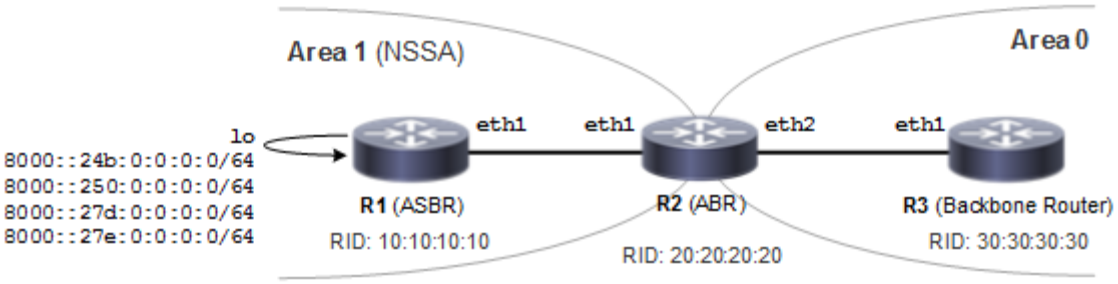

Figure 3-54: Using the summary-address Option

Configuration

R1

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 address 1000::1/64`	Configure ipv6 address for interface eth1
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config-if)#interface lo`	Enter interface mode for loopback
`(config-if)#ipv6 address 8000::24b:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::250:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::27d:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#ipv6 address 8000::27e:0:0:0:0/64`	Assign IPv6 address to loopback interface
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 10.10.10.10`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa`	Configure the area as NSSA.
`(config-router)#redistribute connected`	Redistribute the configured loopback network into OSPFv3 NSSA. Note: Connected networks can be redistributed by setting the metric and metric type.
`(config-router)#summary-address 8000::/48 all-tag 10`	Summarize the address range and tag the summarized route
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

R2

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 address 1000::2/64`	Configure ipv6 address for interface eth1
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode for `eth2`
`(config-if)#ipv6 address 2000::1/64`	Configure ipv6 address for interface eth2
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 20.20.20.20`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa`	Configure the Router in NSSA
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

R3

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 address 2000::2/64`	Configure ipv6 address for interface eth1
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 30.30.30.30`	Configure the router ID to use on this instance (100)
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

In the configurations above, you can suppress the external route summarization by NSSA-ASBR by specifying the `not-advertise` parameter as shown below:

```
(config-router)#summary-address 8000::/48 not-advertise
```

Also, connected networks can be redistributed by setting the metric and metric type as shown below:

```
(config-router)#redistribute connected metric 20 metric-type 1
```

Validation 1

The output below shows the summarized route generated by NSSA-ASBR (R1) with a tag. The output has the LS Type as NSSA-external-LSA with advertising router identifier (10.10.10.10) of the NSSA-ASBR (R1). Also, check the Prefix which is summarized route and external route tag as configured.

```
R1#sh ipv6 ospf database nssa-external

            OSPFv3 Router with ID (10.10.10.10) (Process 100)

                NSSA-external-LSA (Area 0.0.0.1)
```

```
  LS age: 90
  LS Type: NSSA-external-LSA
  Link State ID: 0.0.0.11
  Advertising Router: 10.10.10.10
  LS Seq Number: 0x80000003
  Checksum: 0x69B3
  Length: 40
    Metric Type: 2 (Larger than any link state path)
    Metric: 21
    Prefix: 8000::/48
    Prefix Options: 8 (P|-|-|-)
    External Route Tag: 10
```

Validation 2

The output below on the NSSA-ABR that is translating Type-7 LSAs to Type-5 LSAs shows summarized address in Type-7 and Type-5 LSA. Check for the same prefix, external route tag in both Type7 and Type-5 LSA.

```
R2#sh ipv6 ospf database nssa-external

            OSPFv3 Router with ID (20.20.20.20) (Process 100)

                NSSA-external-LSA (Area 0.0.0.1)

  LS age: 241
  LS Type: NSSA-external-LSA
  Link State ID: 0.0.0.11
  Advertising Router: 10.10.10.10
  LS Seq Number: 0x80000003
  Checksum: 0x69B3
  Length: 40
    Metric Type: 2 (Larger than any link state path)
    Metric: 21
    Prefix: 8000::/48
    Prefix Options: 8 (P|-|-|-)
    External Route Tag: 10

R2#sh ipv6 ospf database external

            OSPFv3 Router with ID (20.20.20.20) (Process 100)

                AS-external-LSA

  LS age: 245
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.3
  Advertising Router: 20.20.20.20
  LS Seq Number: 0x80000003
  Checksum: 0x8660
  Length: 40
    Metric Type: 2 (Larger than any link state path)
    Metric: 21
    Prefix: 8000::/48
    Prefix Options: 0 (-|-|-|-)
    External Route Tag: 10
```

Validation 3

The output below on the backbone router shows the summarized address in the translated Type-5 LSA. The prefix and external route tag are the same as the summarized Type-7 LSA originated by R1.

```
R3#sh ipv6 ospf database external

            OSPFv3 Router with ID (30.30.30.30) (Process 100)

                AS-external-LSA

  LS age: 409
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.3
  Advertising Router: 20.20.20.20
  LS Seq Number: 0x80000003
  Checksum: 0x8660
  Length: 40
    Metric Type: 2 (Larger than any link state path)
    Metric: 21
    Prefix: 8000::/48
    Prefix Options: 0 (-|-|-|-)
    External Route Tag: 10
```

NSSA with the Translator Role Option

Type-7 to Type-5 translation is done by an NSSA-ABR. If an NSSA has multiple NSSA-ABRs, only one will perform the translation. The NSSA-ABR translator role options are:

- Candidate (default)
- Always

In the topology in Figure 3-55:

- R1 is NSSA-ASBR
- R2 and R3 are NSSA-ABRs
- R4 is a backbone router

In this example, the NSSA translator role `candidate` is configured on both NSSA-ABRs (R2 and R3). The Type-7 to Type-5 translation is done by the router with the higher router identifier (R3).

Topology

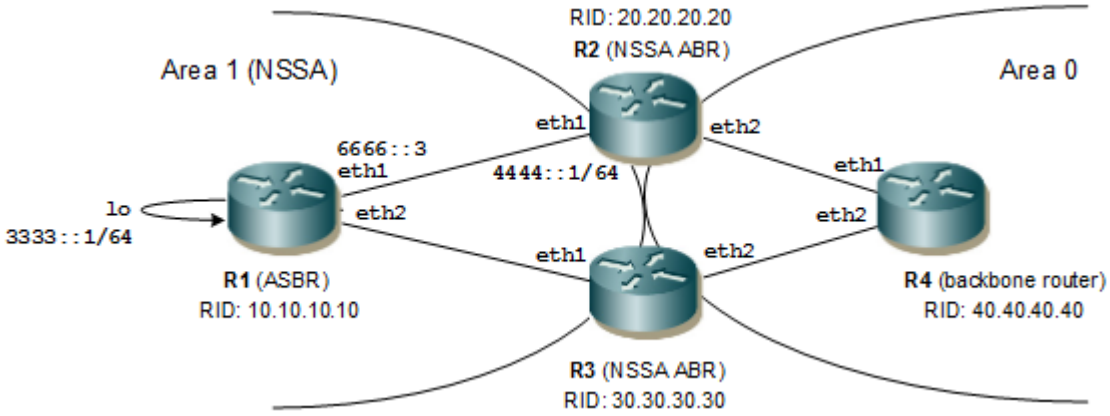

Figure 3-55: Using the translator-role Option

Configuration

When one NSSA-ABR is configured with the translator role as `always` and the other as `candidate`, then translation is done by the router configured as `always`. In this scenario, the translation can be biased by setting the translator role to `always` on the router that has the lower router identifier.

R1

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode for `eth2`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config-if)#interface lo`	Enter interface mode for Loopback
`(config-if)#ipv6 address 3333::1/64`	Assign IPv6 address to loopback interface
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 10.10.10.10`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa`	Configure the area as NSSA.
`(config-router)#redistribute static`	Redistribute the static route configured into the OSPF NSSA
`(config-router)#redistribute connected`	Redistribute the connected network into OSPF NSSA
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode
`(config)#ipv6 route 4444::1:0:0:0:0/64 6666::3`	Configure the static route with the nexthop address as R2's `eth1` IPv6 address
`(config)#exit`	Exit interface mode.

R2

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode for `eth2`
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 20.20.20.20`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa translator-role candidate`	Configure the NSSA-ABR with the translator role candidate.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

R3

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 1`	Configure interface in an area assigned with the area ID (1).
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode for `eth2`
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 30.30.30.30`	Configure the router ID to use on this instance (100)
`(config-router)#area 1 nssa translator-role candidate`	Configure the NSSA-ABR with the translator role candidate.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

R4

`(config)#interface eth1`	Enter interface mode for `eth1`.
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config)#interface eth2`	Enter interface mode for `eth2`
`(config-if)#ipv6 router ospf tag 100 area 0`	Configure interface in backbone area (0)
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 100`	Configure the routing process and specify the tag (`100`).
`(config-router)#router-id 40.40.40.40`	Configure the router ID to use on this instance (100)
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit interface mode

The command to configure the NSSA-Translator role as `always` is:

```
(config-router)#area 1 nssa translator-role always
```

The NSSA-ABR can continue to perform translation after its services are no longer required for the stability interval which is set using the command below on the NSSA-ABR.

```
(config-router)#area 1 nssa stability-interval 7777
```

Validation 1

The translation is done by the NSSA-ABR with the higher router identifier. In the output below, check the router identifier of the NSSA-ABR. Also, check the router which is elected and the router which is disabled.

```
R2#sh ipv6 ospf
 Routing Process "OSPFv3 (100)" with ID 20.20.20.20
 Process uptime is 21 minutes
 SPF schedule delay min 0.500 secs, SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incoming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 4. Checksum Sum 0x1F816
 Number of AS-Scoped Unknown LSA 0
 Number of LSA originated 28
 Number of LSA received 58
 Number of areas in this router is 2
    Area BACKBONE(0)
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 7 times
        Number of LSA 19.  Checksum Sum 0x7454D
        Number of Unknown LSA 0
    Area 0.0.0.1 (NSSA)
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 14 times
        Number of LSA 19.  Checksum Sum 0xA4D18
        Number of Unknown LSA 0
        NSSA Translator State is disabled
R3#sh ipv6 ospf
 Routing Process "OSPFv3 (100)" with ID 30.30.30.30
 Process uptime is 19 minutes
 SPF schedule delay min 0.500 secs, SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incomming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 4. Checksum Sum 0x1F816
 Number of AS-Scoped Unknown LSA 0
 Number of LSA originated 31
 Number of LSA received 69
 Number of areas in this router is 2
    Area BACKBONE(0)
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 15 times
        Number of LSA 19.  Checksum Sum 0x7454D
        Number of Unknown LSA 0
    Area 0.0.0.1 (NSSA)
        Number of interfaces in this area is 1(1)
        SPF algorithm executed 10 times
        Number of LSA 19.  Checksum Sum 0xA4D18
```

```
        Number of Unknown LSA 0
        NSSA Translator State is elected
```

Validation 2

The translated Type-5 LSA in R4 in area 0 has the advertising router identifier of R3. In the output below, the LS Type is AS-External-LSA and the advertising router has the higher router identifier.

```
R4#sh ipv6 ospf database external

            OSPFv3 Router with ID (40.40.40.40) (Process 100)

                AS-external-LSA

  LS age: 885
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.7
  Advertising Router: 30.30.30.30
  LS Seq Number: 0x80000001
  Checksum: 0xD3FE
  Length: 40
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 3333::/64
    Prefix Options: 0 (-|-|-|-)
    External Route Tag: 0


  LS age: 18
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.8
  Advertising Router: 30.30.30.30
  LS Seq Number: 0x80000003
  Checksum: 0x7457
  Length: 56
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 4444::/64
    Prefix Options: 0 (-|-|-|-)
    Forwarding Address: 6666::3
    External Route Tag: 0
```

Link LSA Suppression

If link LSA suppression is enabled and the interface type is not broadcast or NBMA, the router will not originate a link-LSA for the link. This implies that other routers on that link will determine the router's next hop address using a mechanism other than the link LSA.

Topology

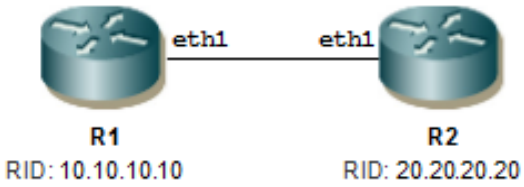

Figure 3-56: LSA Suppression

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process.
`(config-router)#router-id 10.10.10.10`	Configure the router ID to use on this instance.
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 router ospf tag 100-ABC area 1`	Configure interface in an area assigned with the area ID (1). The tag uniquely identifies the routing process.
`(config-if)#ipv6 ospf network point-to-point`	Configure the OSPF interface network type as point to point
`(config-if)#ipv6 ospf link-lsa-suppression enable`	Enable the link LSA suppression mechanism
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

R2

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process.
`(config-router)#router-id 20.20.20.20`	Configure the router ID to use on this instance.
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 router ospf tag 100-ABC area 1`	Configure interface in an area assigned with the area ID (1). The tag uniquely identifies the routing process.
`(config-if)#ipv6 ospf network point-to-point`	Configure the OSPF interface network type as point to point
`(config-if)#ipv6 ospf link-lsa-suppression enable`	Enable the link LSA Suppression Mechanism
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

Note: This is not applicable for broadcast and NBMA networks.

Validation 1

Verify that adjacency has been established.

```
R1#sh ipv6 ospf neighbor
OSPFv3 Process (100)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
20.20.20.20       1   Full/ -         00:00:37    eth1       0
```

Validation 2

Verify that R1 should not have the Link LSA in the Link state database.

Note: The output below is captured after link lsa suppression enabled which has not Link LSA in the LSDB.

```
R1#sh ipv6 ospf database

            OSPFv3 Router with ID (10.10.10.10) (Process 100-ABC)

                Router-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age  Seq#       CkSum    Link
0.0.0.0         10.10.10.10       15 0x80000004 0x3264      1
0.0.0.0         20.20.20.20       15 0x80000002 0xdbba      1

                Intra-Area-Prefix-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age  Seq#       CkSum  Prefix  Reference
0.0.0.12        10.10.10.10       14 0x80000004 0xaab4      1  Router-LSA
0.0.0.13        20.20.20.20       15 0x80000002 0x8f7f      1  Router-LSA

                Intra-Area-Te-LSA (Area 0.0.0.1)

Link State ID   ADV Router      Age  Seq#       CkSum
0.0.0.4         10.10.10.10       15 0x80000004 0xa326
0.0.0.3         20.20.20.20       15 0x80000002 0xffec
R1#sh ipv6 ospf database link

            OSPFv3 Router with ID (10.10.10.10) (Process 100-ABC)
```

Originate Type-7 LSAs and Translate to Type-5

Figure 3-57 shows the configuration to originate Type-7 LSAs and translate them into Type-5 LSAs. R3 is an NSSA-ASBR that originates Type-7 LSAs into the NSSA which are converted to Type-5 LSAs by R2 which is an NSSA-ABR. R1 is a backbone router.

Topology

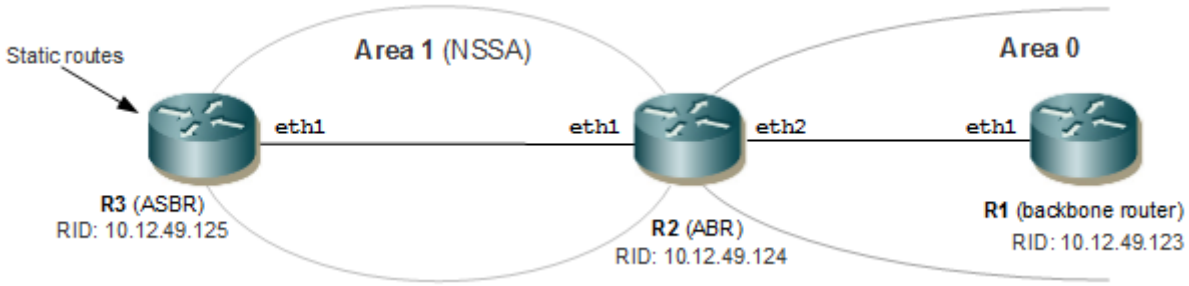

Figure 3-57: Originate Type-7 LSAs and Translate to Type-5 under Address Family IPv4

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process
`(config-router)#router-id 10.12.49.123`	Configure the router ID to use on this tag
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 router ospf area 0 tag 100-ABC instance-id 31`	Configure the interface in an area assigned with the area ID (0) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv6 address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

R2

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process.
`(config-router)#ip route 15.15.15.0/24 null`	Configure the static route with the nexthop address set to null
`(config-router)#router-id 10.12.49.125`	Configure the router ID to use for this process
`(config-router)#area 1 nssa`	Configure the area 1 as NSSA.
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 router ospf area 0 tag 100-ABC instance-id 31`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 64-95 for the IPv4 address family.
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode

© 2025 IP Infusion Inc. Proprietary

`(config-if)#ipv6 router ospf area 1 tag 100-ABC instance-id 30`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 0-30 for the IPv6 address family.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

R3

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process.
`(config-router)#router-id 10.12.49.125`	Configure the router ID to use for this process
`(config-router)#area 1 nssa`	Configure the area 1 as NSSA.
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ipv6 router ospf area 1 tag 100-ABC instance-id 30`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv6 address family.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation 1

Verify that adjacency has been established with the configured instance identifier.

```
R2#sh ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface   Instance ID
10.12.49.123      1   Full/DR         00:00:31       eth1           31
10.12.49.125      1   Full/Backup     00:00:38    eth2          30
```

Validation 2

Verify that R3 has generated a Type-7 LSA and that the ABR R2 has External LSA Type 5 in its Database.

R3

```
R3#show ipv6 ospf database nssa-external

            OSPFv3 Router with ID (10.12.49.125) (Process 100-ABC)

                NSSA-external-LSA (Area 0.0.0.1)

  LS age: 139
  LS Type: NSSA-external-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0xAB34
  Length: 48
```

```
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 15.15.15.0/24
    Prefix Options: 8 (P|-|-|-)
    Forwarding Address: 22.1.1.2
    External Route Tag: 0


R3#show ipv6 ospf database external

            OSPFv3 Router with ID (10.12.49.125) (Process 100-ABC)

R3#
```

Validation 3

R2

```
R2#show ipv6 ospf database nssa-external

            OSPFv3 Router with ID (10.12.49.124) (Process 100-ABC)

                NSSA-external-LSA (Area 0.0.0.1)

  LS age: 105
  LS Type: NSSA-external-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0xAB34
  Length: 48
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 15.15.15.0/24
    Prefix Options: 8 (P|-|-|-)
    Forwarding Address: 22.1.1.2
    External Route Tag: 0

R2#

R2#show ipv6 ospf database external

            OSPFv3 Router with ID (10.12.49.124) (Process 100-ABC)

                AS-external-LSA

  LS age: 706
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.124
  LS Seq Number: 0x80000001
  Checksum: 0xAB1F
  Length: 48
```

```
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 15.15.15.0/24
    Prefix Options: 0 (-|-|-|-)
    Forwarding Address: 22.1.1.2


R2#

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.0.0.0/8 is directly connected, eth0, 15:43:05
C            11.1.1.0/24 is directly connected, eth1, 14:54:49
O N2         15.15.15.0/24 [110/20] via 22.1.1.2, eth2, 12:09:25
C            22.1.1.0/24 is directly connected, eth2, 12:22:45
C            127.0.0.0/8 is directly connected, lo, 15:43:05

Gateway of last resort is not set
R2#
```

Validation 3

Verify that FIB of backbone router has External Route as “O E2”.

```
R1#
Verify that FIB of backbone router R1 has External Route as "O E2".
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.0.0.0/8 is directly connected, eth0, 03:34:25
C            11.1.1.0/24 is directly connected, eth1, 02:46:18
O E2         15.15.15.0/24 [110/20] via 11.1.1.2, eth1, 00:00:36
O IA         22.1.1.0/24 [110/2] via 11.1.1.2, eth1, 00:05:01
C            127.0.0.0/8 is directly connected, lo, 03:34:25

Gateway of last resort is not set
```

Summarize Inter-Area and External Routes

Figure 3-58 shows the configuration to enable inter-area and external route summarization. The IPv4 address family is enabled on R1. R2 summarizes the internal OSPF routes which R3 redistributes.

Topology

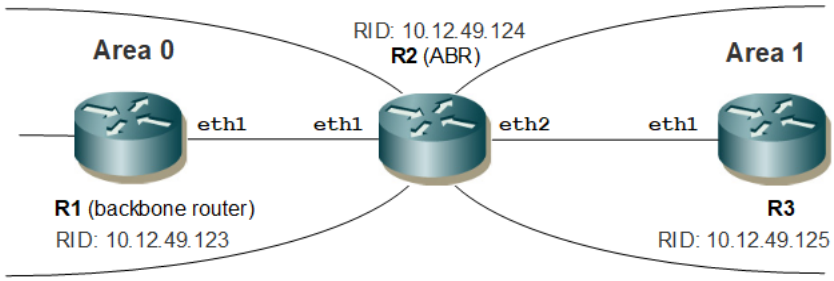

Figure 3-58: Enabling Intra-Area and External Route Summarization

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (`100-ABC`) which uniquely identifies the routing process
`(config-router)#router-id 10.12.49.123`	Configure the router ID to use for this process.
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 10.10.10.1/24`	Specify IP address for interface eth1
`(config-if)#ipv6 router ospf area 0 tag 100-ABC instance-id 31`	Configure the interface in an area assigned with the area ID (0) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv6 address family.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (100-ABC) which uniquely identifies the routing process
`(config-router)#router-id 10.12.49.124`	Configure the router ID to use for this process
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 10.10.10.2/24`	Specify an IP address for the interface
`(config-if)#ipv6 router ospf area 0 tag 100-ABC instance-id 31`	Configure the interface in an area assigned with the area ID (0) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv4 address family.
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 20.20.20.1/24`	Specify an IP address for the interface
`(config-if)#ipv6 router ospf area 1 tag 100-ABC instance-id 30`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv6 address family.
`(config-if)#exit`	Exit interface mode
`(config)#interface lo`	Enter interface mode
`(config-if)#ipv6 router ospf area 1 tag 101 instance-id 30`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv6 address family.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

R3

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 100-ABC`	Configure the routing process and specify the tag (100-ABC) which uniquely identifies the routing process
`(config-router)#router-id 10.12.49.125`	Configure the router ID to use for this process
`(config-router)#exit`	Exit OSPF router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 20.20.20.2/24`	Specify an IP address for the interface
`(config-if)#ipv6 router ospf area 1 tag 100-ABC instance-id 30`	Configure the interface in an area assigned with the area ID (1) which uniquely identifies the routing process and the instance identifier which is 0-31 for the IPv4 address family.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode

Validation

Validation 1: Verify that adjacency has been established with the configured instance identifier.

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (100-ABC)
Neighbor ID     Pri   State              Dead Time   Interface  Instance ID
10.12.49.123      1   Full/Backup     00:00:38       eth1       31
10.12.49.125      1   Full/DR            00:00:38       eth2       30
```

Validation 2: Verify that a single summarized OSPF IA route and a single summarized external route is available in FIB of R1

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.0.0.0/8 is directly connected, eth0, 00:45:18
C            11.1.1.0/24 is directly connected, eth1, 00:40:01
O IA         22.1.1.0/24 [110/2] via 11.1.1.2, eth1, 00:37:57
O IA         100.1.1.0/24 [110/2] via 11.1.1.2, eth1, 00:24:59
C            127.0.0.0/8 is directly connected, lo, 00:45:18
O E2         200.1.1.0/24 [110/20] via 11.1.1.2, eth1, 00:00:54

Gateway of last resort is not set
R1#
```

Validation 3: Verify that the Inter-Area Prefix LSA and External LSA in OSPFv3 database of R1 consists of just a single prefix 100.1.1.0/24 and 200.1.1.0/24 respectively

```
R1#show ipv6 ospf database inter-prefix

            OSPFv3 Router with ID (10.12.49.123) (Process 100-ABC)

                Inter-Area-Prefix-LSA (Area 0.0.0.0)

  LS age: 771
  LS Type: Inter-Area-Prefix-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.124
  LS Seq Number: 0x80000002
  Checksum: 0x60E3
  Length: 32
    Metric: 1
    Prefix: 22.1.1.0/24
    Prefix Options: 0


  LS age: 21
  LS Type: Inter-Area-Prefix-LSA
```

```
  Link State ID: 0.0.0.2
  Advertising Router: 10.12.49.124
  LS Seq Number: 0x80000008
  Checksum: 0x489D
  Length: 32
    Metric: 1
    Prefix: 127.0.0.0/8
    Prefix Options: 0


  LS age: 1795
  LS Type: Inter-Area-Prefix-LSA
  Link State ID: 0.0.0.5
  Advertising Router: 10.12.49.124
  LS Seq Number: 0x80000001
  Checksum: 0x975B
  Length: 32
    Metric: 1
    Prefix: 100.1.1.0/24
    Prefix Options: 0


R1#


R1#show ipv6 ospf database external

            OSPFv3 Router with ID (10.12.49.123) (Process 100-ABC)

                AS-external-LSA

  LS age: 390
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0xCE5A
  Length: 32
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 10.0.0.0/8
    Prefix Options: 0 (-|-|-|-)


  LS age: 364
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.4
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0x6CE8
  Length: 32
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 200.1.1.0/24
    Prefix Options: 0 (-|-|-|-)
```

```
R1#
```

Validation 4: Verify that a single summarized external route is present in the ABR R2

```
R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.0.0.0/8 is directly connected, eth0, 00:55:15
C            11.1.1.0/24 is directly connected, eth1, 00:49:17
C            22.1.1.0/24 is directly connected, eth2, 00:48:18
O            100.1.1.0/24 [110/0] is a summary, Null, 00:35:05
C            100.1.1.100/32 is directly connected, lo, 00:35:40
C            100.1.1.110/32 is directly connected, lo, 00:35:36
C            127.0.0.0/8 is directly connected, lo, 00:55:15
O E2         200.1.1.0/24 [110/20] via 22.1.1.2, eth2, 00:11:00

Gateway of last resort is not set
R2#
```

Validation 5: Verify that the Type 5 LSA in the ABR's Link State Data Base consistes of just a single prefix 200.1.1.0/24

```
R2#show ipv6 ospf database external

            OSPFv3 Router with ID (10.12.49.124) (Process 100-ABC)

                AS-external-LSA

  LS age: 774
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.1
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0xCE5A
  Length: 32
    Metric Type: 2 (Larger than any link state path)
    Metric: 20
    Prefix: 10.0.0.0/8
    Prefix Options: 0 (-|-|-|-)


  LS age: 748
  LS Type: AS-External-LSA
  Link State ID: 0.0.0.4
  Advertising Router: 10.12.49.125
  LS Seq Number: 0x80000001
  Checksum: 0x6CE8
  Length: 32
```

```
Metric Type: 2 (Larger than any link state path)
Metric: 20
Prefix: 200.1.1.0/24
Prefix Options: 0 (-|-|-|-)
```

Distribute List

To filter the routes that Open Shortest Path First Version 3 (OSPFv3) installs in the Routing Information Base (RIB), use the `distribute-list in` command in an appropriate configuration mode.

To filter the routes redistributed into Open Shortest Path First Version 3 (OSPFv3) from other routing protocols, use the `distribute-list out` command in an appropriate configuration mode.

Topology

Figure 3-59 shows the configuration to illustrate the distribute-list support for OSPFv3

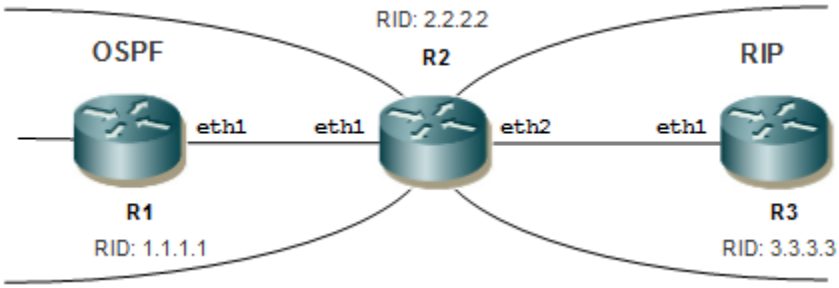

Figure 3-59: Basic Topology for Distribute-list

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 address 2000::1/64`	Configure the IPv6 address of the interface.
`(config-if)#ipv6 router ospf area 0 tag proc1`	Configure the interface in an area assigned with the area ID (0) which uniquely identifies the routing process
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)# ipv6 address 1111::1/128`	Configure the IPv6 address of the interface.
`(config-if)# ipv6 address 2222::2/128`	Configure the IPv6 address of the interface.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf proc1`	Configure the routing process
`(config-router)#router-id 1.1.1.1`	Configure router-id to uniquely identify the router
`(config-router)#redistribute connected`	Redistribute connected routes into ospfv3
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#end`	Exit router mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 address 2000::50/64`	Configure the IPv6 address of the interface
`(config-if)#ipv6 router ospf area 0 tag proc1`	Configure the interface in an area assigned with the area ID (0) which uniquely identifies the routing process
`(config-if)#exit`	Exit interface mode
`(config)#interface eth2`	Enter interface mode
`(config-if)#ipv6 address 4000::50/64`	Configure the IPv6 address of the interface.
`(config-if)#exit`	Exit interface mode
`(config-if)#ipv6 router rip`	Configure rip instance under interface
`(config-if)#exit`	Exit interface mode
`(config)#router ipv6 rip`	Configure the rip routing process
`(config-router)#neighbor fe80::5054:ff:fe85:19bc eth2`	Configure RIP neighbor peer
`(config-router)#exit`	Exit router mode.
`(config)#ipv6 access-list 1`	Configure ipv6 access list
`(config-ipv6-acl)# permit any 7777::/64 any`	Configure ipv6 access-list to permit 7777::/64 and deny 8888::/64
`(config-ipv6-acl)#exit`	Exit ipv6 access-list mode
`(config)#ipv6 access-list 2`	Configure ipv6 access-list

© 2025 IP Infusion Inc. Proprietary

(config-ipv6-acl)#permit any 1111::1/128 any	Configure ipv6 access-list to permit 1111::1/128 and deny 2222::2/128
(config-ipv6-acl)#exit	Exit ipv6 access-list mode
(config)#router ipv6 ospf proc1	Configure the ospfv3 routing process
(config-router)#router-id 2.2.2.2	Configure router-id to uniquely identify the router
(config-router)#redistribute rip	Redistribute rip routes
(config-router)#distribute-list 1 out rip	Configure distribute list to allow only the permitted routes redistributed from RIP
(config-router)#distribute-list 2 in	Configure distribute list to allow the installation of only the permitted OSPFv3 routes in RIB
(config-router)#exit	Exit router mode
(config)#ipv6 access-list 1	Enter access-list mode
(config-ipv6-acl)#permit any 8888::/64 any	Configure the ipv6 access-list to permit 8888::/64 alongwith 7777::/64
(config-ipv6-acl)#exit	Exit access-list mode
(config)#ipv6 access-list 2	Enter access-list mode
(config-ipv6-acl)#permit any 2222::2/128 any	Configure the ipv6 access-list to permit 2222::2/128 alongwith 1111::1/128
(config-ipv6-acl)#exit	Exit access-list mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#exit	Exit configure mode

R3

#configure terminal	Enter configure mode
(config)#interface eth1	Enter interface mode
(config-if)#ipv6 address 4000::51/64	Configure the IPv6 address of the interface.
(config-if)#ipv6 router rip	Configure rip instance under interface
(config-if)#exit	Exit interface mode
(config)#router ipv6 rip	Configure the rip routing process
(config-router)#neighbor fe80::5054:ff:fec6:69f eth1	Configure rip neighbor peer
(config-router)#exit	Exit router mode
(config)#ipv6 route 7777::/64 eth2	Configure static route
(config)#ipv6 route 8888::/64 eth3	Configure static route
(config)#router ipv6 rip	Configure the rip routing process
(config-router)#redistribute static	Redistribute configured static routes
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode

Validation 1

Verify OSPF neighborship is up between R1and R2

R2

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (Proc1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:00:38    eth1       0
```

Validation 2

Check if permitted route 7777::/64 is present in R1's routing table and denied route 8888::/64 is not present.

R1

```
R1#show ipv6 ospf route
OSPFv3 Process (Proc1)
Codes: C - connected, D - Discard, 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
   Destination     Metric           Next-hop
C  2000::/64         1             directly connected, eth1, Area 0.0.0.0
E2 7777::/64        1/20   via    fe80::5054:ff:fe1e:269d, eth1
```

Validation 3

Check both the routes 7777::/64 and 8888::/64 are present after 8888::/64 is permitted

R1

```
rtr1#show ipv6 ospf route
OSPFv3 Process (Proc1)
Codes: C - connected, D - Discard, 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

   Destination    Metric    Next-hop
C  2000::/64         1      directly connected, eth1, Area 0.0.0.0
E2 7777::/64        1/20   via fe80::5054:ff:fe1e:269d, eth1
E2 8888::/64        1/20   via fe80::5054:ff:fe1e:269d, eth1
```

Validation 4

Check if permitted route 1111::1/128 is present in R2's routing table and denied route 2222::2/128 is not present.

R1

```
R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
```

© 2025 IP Infusion Inc. Proprietary

```
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:43:35
O E2    1111::1/128 [110/20] via fe80::5054:ff:fe0c:40ed, eth1, 00:01:17
C       2000::/64 via ::, eth1, 00:36:33
C       4000::/64 via ::, eth2, 00:36:19
R       7777::/64 [120/2] via fe80::5054:ff:fe96:a3f9, eth2, 00:21:57
R       8888::/64 [120/2] via fe80::5054:ff:fe96:a3f9, eth2, 00:21:57
C       fe80::/64 via ::, eth9, 00:43:35
R2#
```

Validation 5

Check both the routes 1111::1/128 and 2222::2/128 are present after 2222::2/128 is permitted.

R1

```
R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:54:52
O E2    1111::1/128 [110/20] via fe80::5054:ff:fe0c:40ed, eth1, 00:12:34
C       2000::/64 via ::, eth1, 00:47:50
O E2    2222::2/128 [110/20] via fe80::5054:ff:fe0c:40ed, eth1, 00:00:02
C       4000::/64 via ::, eth2, 00:47:36
R       7777::/64 [120/2] via fe80::5054:ff:fe96:a3f9, eth2, 00:33:14
R       8888::/64 [120/2] via fe80::5054:ff:fe96:a3f9, eth2, 00:33:14
C       fe80::/64 via ::, eth9, 00:54:52
```

OSPFv3 Authentication

This example shows the configuration required for enabling OSPFv3 authentication with IPSEC on an OSPFv3-enabled interface. R1 and R2 are two routers in Area 0 connecting to the network 2000::/64.

Note: You must explicitly specify a Router ID for the OSPFv3 process to be activated.

Topology

Figure 3-60: OSPFv3 Authentication

R1

`#configure terminal`	Enter configure mode.
`(config)#crypto ipsec transform-set tset1 esp-auth esp-md5 esp-enc esp-3des`	Create a transform-set with the ESP Authentication and encryption parameters
`(config)#crypto map map1 ipsec-manual`	Create manual key association to the crypto-map
`(config-crypto)#sequence 100`	Create sequence in crypto-map
`(config-crypto-seq)# set transform-set tset1`	Attach transform set to cryptomap
`(config-crypto-seq)# set peer fe80::923c:b3ff:fe82:8d88 spi 2`	Set IPv6 peer (OSPFv3 neighbor) and SPI value
`(config-crypto-seq)# set session-key outbound esp 1 cipher 66546A576E5A72346A586E3272357538782F413F4428472B authenticator 3777217A25432A46763979244226452948404D6351655468`	Set the outbound session-key with cipher and auth values
`(config-crypto-seq)#exit`	Exit sequence mode in crypto
`(config-crypto)#exit`	Exit crypto mode
`(config)#router ipv6 ospf 1`	Create an OSPFv3 routing instance.
`(config-router)#router-id 1.1.1.1`	Specify a Router ID for the OSPFv3 routing process.
`(config-router)#exit`	Exit OSPF router mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router ospf area 0 tag 1`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`(config-if)# ipv6 ospf authentication cryptomap map1`	Configure OSPFv3 authentication over the OSPFv3 interface by applying the configured cryptomap
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#crypto ipsec transform-set tset1 esp-auth esp-md5 esp-enc esp-3des`	Create a transform-set with the ESP Authentication and encryption parameters
`(config)#crypto map map1 ipsec-manual`	Create manual key association to the crypto-map
`(config-crypto)#sequence 100`	Create sequence in crypto-map
`(config-crypto-seq)# set transform-set tset1`	Attach transform set to cryptomap
`(config-crypto-seq)# set peer fe80::fa8e:a1ff:fe0b:dd9a spi 1`	Set IPv6 peer (OSPFv3 neighbor) and SPI value
`(config-crypto-seq)# set session-key outbound esp 2 cipher 66546A576E5A72346A586E3272357538782F413F4428472B authenticator 3777217A25432A46763979244226452948404D6351655468`	Set the outbound session-key with cipher and auth values
`(config-crypto-seq)#exit`	Exit sequence mode in crypto
`(config-crypto)#exit`	Exit crypto mode
`(config)#router ipv6 ospf 1`	Create an OSPFv3 routing instance.

(config-router)#router-id 2.2.2.2	Specify a Router ID for the OSPFv3 routing process.
(config-router)#exit	Exit OSPF router mode.
(config)#interface eth2	Enter interface mode.
(config-if)#ipv6 router ospf area 0 tag 1	Enable OSPFv3 routing on an interface, and assign the Area ID (0).
(config-if)# ipv6 ospf authentication cryptomap map1	Configure OSPFv3 authentication over the OSPFv3 interface by applying the configured cryptomap
(config-if)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
R1#show running-config ipsec
!
crypto ipsec transform-set tset1 esp-auth esp-md5 esp-enc esp-3des

crypto map map1 ipsec-manual
 sequence 100
  set transform-set tset1
  set peer fe80::fa8e:a1ff:fe0b:dd9a spi 2
  set session-key outbound esp 1 cipher
66546A576E5A72346A586E3272357538782F413F4428472B authenticator
3777217A25432A467639792442264529
!
R1#
R1#show running-config ospfv3
!
router ipv6 ospf 1
 router-id 1.1.1.1
!
interface eth1
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 ospf authentication cryptomap map1
!

R1#show ipv6 ospf neighbor

Total number of full neighbors: 1
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR         00:00:12    eth1        0

R1#show ipv6 ospf neighbor detail
 Neighbor 2.2.2.2, interface address fe80::923c:b3ff:fe82:8d88
    In the area 0.0.0.0 via interface ce0
    Neighbor priority is 1, State is Full, 6 state changes
    DR is 2.2.2.2 BDR is 1.1.1.1
    Options is 0x000113 (AF|*|*|-|R|-|-|E|V6)
    Dead timer due in 00:00:34
```

```
    Database Summary List 0
    Link State Request List 0
    Link State Retransmission List 0
    Bidirectional Forwarding Detection is enabled

R1#show crypto ipsec sadb
SRC: fe80::923c:b3ff:fe82:8d88      DST:fe80::fa8e:a1ff:fe0b:dd9a
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used
Used bytes=0 Used Count=0

SRC: fe80::923c:b3ff:fe82:8d88     DST:ff02::6
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used
Used bytes=0 Used Count=0

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::6
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used
Used bytes=0 Used Count=0

SRC: fe80::923c:b3ff:fe82:8d88     DST:ff02::5
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Wed Jun 22 03:32:20 2022
Used bytes=0 Used Count=0

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::5
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Wed Jun 22 03:32:21 2022
Used bytes=908 Used Count=25

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:2000::2
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used
Used bytes=0 Used Count=0

R1#show crypto ipsec spdb
SRC:2000::2     DST:ff02::6
Policy Type=ipsec Dir=in
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used

SRC:2000::2     DST:ff02::5
Policy Type=ipsec Dir=in
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used

SRC:2000::2     DST:fe80::fa8e:a1ff:fe0b:dd9a
Policy Type=ipsec Dir=in
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::6
```

```
Policy Type=ipsec Dir=out
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::5
Policy Type=ipsec Dir=out
Added at=Wed Jun 22 03:32:11 2022 First Used at=Wed Jun 22 03:35:02 2022

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:2000::2
Policy Type=ipsec Dir=out
Added at=Wed Jun 22 03:32:11 2022 First Used at=Never Used

root@R1:~# tcpdump -i eth1
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth1, link-type EN10MB (Ethernet), capture size 262144 bytes
03:33:20.393469 IP6 fe80::923c:b3ff:fe82:8d88 > ff02::5: ESP(spi=0x00000002,seq=0x8),
length 68
03:33:21.174899 IP6 fe80::fa8e:a1ff:fe0b:dd9a > ff02::5: ESP(spi=0x00000001,seq=0x7),
length 68
03:33:30.394602 IP6 fe80::923c:b3ff:fe82:8d88 > ff02::5: ESP(spi=0x00000002,seq=0x9),
length 68
03:33:32.175157 IP6 fe80::fa8e:a1ff:fe0b:dd9a > ff02::5: ESP(spi=0x00000001,seq=0x8),
length 68
```

R2

```
R2#show running-config ipsec
!
crypto ipsec transform-set tset1 esp-auth esp-md5 esp-enc esp-3des

crypto map map1 ipsec-manual
 sequence 100
  set transform-set tset1
  set peer fe80::fa8e:a1ff:fe0b:dd9a spi 1
  set session-key outbound esp 2 cipher
462D4A614E6452675166546A576E5A723475377821412544 authenticator
3777217A25432A467639792442264529
!

R2#show running-config ospfv3
!
router ipv6 ospf 1
 router-id 2.2.2.2
 bfd all-interfaces
!
interface ce1
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 ospf authentication cryptomap map1
!

R2#show ipv6 ospf neighbor
Total number of full neighbors: 1
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
```

```
1.1.1.1           1   Full/Backup     00:00:31    ce1        0

R2#show ipv6 ospf neighbor detail
 Neighbor 1.1.1.1, interface address fe80::fa8e:a1ff:fe0b:dd9a
    In the area 0.0.0.0 via interface ce1
    Neighbor priority is 1, State is Full, 6 state changes
    DR is 2.2.2.2 BDR is 1.1.1.1
    Options is 0x000113 (AF|*|*|-|R|-|-|E|V6)
    Dead timer due in 00:00:37
    Database Summary List 0
    Link State Request List 0
    Link State Retransmission List 0
    Bidirectional Forwarding Detection is enabled

R2#show crypto ipsec sadb

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:fe80::923c:b3ff:fe82:8d88
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used
Used bytes=0 Used Count=0

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::6
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used
Used bytes=0 Used Count=0

SRC:fe80::923c:b3ff:fe82:8d88   DST:ff02::6
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used
Used bytes=0 Used Count=0

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::5
SA: spi=0x1  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Mon Feb 18 07:39:14 2019
Used bytes=16072 Used Count=399

SRC:fe80::923c:b3ff:fe82:8d88   DST:ff02::5
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Mon Feb 18 07:39:20 2019
Used bytes=16096 Used Count=399

SRC:fe80::923c:b3ff:fe82:8d88   DST:fe80::fa8e:a1ff:fe0b:dd9a
SA: spi=0x2  sta=MATURE  auth=SHA1HMAC  enc=3DES
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used
Used bytes=0 Used Count=0

R2#show crypto ipsec spdb

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::6
Policy Type=ipsec Dir=in
```

```
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:ff02::5
Policy Type=ipsec Dir=in
Added at=Mon Feb 18 07:39:12 2019 First Used at=Mon Feb 18 08:44:24 2019

SRC:fe80::fa8e:a1ff:fe0b:dd9a   DST:fe80::923c:b3ff:fe82:8d88
Policy Type=ipsec Dir=in
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used

SRC:fe80::923c:b3ff:fe82:8d88   DST:ff02::6
Policy Type=ipsec Dir=out
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used

SRC:fe80::923c:b3ff:fe82:8d88   DST:ff02::5
Policy Type=ipsec Dir=out
Added at=Mon Feb 18 07:39:12 2019 First Used at=Mon Feb 18 08:44:21 2019

SRC:fe80::923c:b3ff:fe82:8d88   DST:fe80::fa8e:a1ff:fe0b:dd9a
Policy Type=ipsec Dir=out
Added at=Mon Feb 18 07:39:12 2019 First Used at=Never Used
```

CHAPTER 4 OSPFv3 Graceful Restart Configuration

The possibility of maintaining a router's data forwarding capability while the router's control software restarts is called graceful restart or non-stop forwarding. After the router restarts and reloads, it must change its OSPF processing until it re-establishes full adjacencies with all its former fully adjacent neighbors. The time period between the restart/reload and re-establishment of adjacencies is called the grace period.

Essentially, the OSPF procedure for graceful restart is as follows:

- The router attempting a graceful restart originates link-local Opaque-LSAs, called Grace-LSAs, announcing its intention to perform a graceful restart within a specified amount of time (grace period).
- During the grace period, neighbors continue to announce the restarting router in their LSAs as if it were fully adjacent (OSPF neighbor state Full), but only if the network topology remains static (the contents of the LSAs in the link-state database that have LS types 1-5 and 7 remain unchanged, and periodic refreshes are allowed).

Note: The Graceful restart configuration is applicable for both OSPFv2 and OSPFv3.

Topology

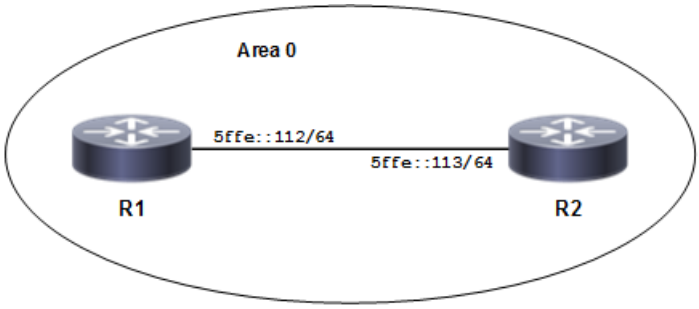

Figure 4-61: OSPFv3 Graceful Restart

Configure R1 for Graceful Restart

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ipv6 address 5ffe::112/64`	Assign the IPv6 address to the interface.
`(config-if)#ipv6 router ospf area 0 tag 1`	Configure the interface for OSPFv3 on area 0.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 1`	Create an OSPFv3 routing instance.
`(config-router)#router-id 1.1.1.1`	Specify a Router ID (`1.1.1.1`) for the OSPFv3 routing process.
`(config-router)#capability restart graceful`	The graceful restart capability is enabled by default. If the user disabled it already, execute this CLI to enable graceful restart capability.
`(config-router)#end`	Exit Configure mode and enter Privileged Exec mode.

`#write`	Save the configuration.
`#restart ipv6 ospf graceful grace-period 100`	Restart OSPFv3 with a 200-second grace period, so that the neighbor maintains adjacency and preserves the routes for 100 seconds.

Configure R2 as Restart Helper

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ipv6 address 5ffe::113/64`	Assign the IPv6 address to the interface.
`(config-if)#ipv6 router ospf area 0 tag 1`	Configure the interface for OSPFv3 on area 0.
`(config-if)#exit`	Exit interface mode.
`(config)#router ipv6 ospf 1`	Create an OSPFv3 routing instance.
`(config-router)#router-id 2.2.2.2`	Specify a Router ID (`2.2.2.2`) for the OSPFv3 routing process.
`(config-router)#exit`	Exit Router mode and enter Configure mode.
`(config)#ipv6 ospf restart helper max-grace-period 1800`	Configure R2 to act as the helper when the grace period is less than 1800.

Remove Capability Restart Configuration from R1

`#configure terminal`	Enter configure mode.
`(config)#router ipv6 ospf 1`	Create an OSPFv3 routing instance.
`(config-router)#router-id 1.1.1.1`	Specify a Router ID (`1.1.1.1`) for the OSPFv3 routing process.
`(config-router)#no capability restart`	Unconfiguring graceful restart capability under router ipv6 ospf 1.
`(config-router)#exit`	Exit Router mode and enter Configure mode.

Remove Helper Configuration from R2

`#configure terminal`	Enter configure mode.
`(config)#ipv6 ospf restart helper never`	Configure R2 to not work as the helper.
`(config)#commit`	Commit the configuration.
`(config)#exit`	Exit Router mode and enter Configure mode.

Validation

R1 After Configuring Graceful Restart

```
R1#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR         00:00:39    eth1       0
```

```
R1#show ipv6 ospf route
OSPFv3 Process (1)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  5ffe::/64                                          1
     directly connected, eth1, Area 0.0.0.0
```

```
rtr1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 02:02:28
C      5ffe::/64 via ::, eth1, 00:23:16
C      fe80::/64 via ::, eth1, 01:25:04
K      ff00::/8 [0/256] via ::, eth0, 01:25:05
```

```
rtr1#show ipv6 ospf database grace

            OSPFv3 Router with ID (1.1.1.1) (Process 1)
```

R1 Before Graceful Restart

```
R1#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR         00:00:37    eth1       0
```

```
rtr1#show ipv6 ospf route
OSPFv3 Process (1)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  5ffe::/64                                          1
     directly connected, eth1, Area 0.0.0.0
```

```
rtr1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 02:02:28
C      5ffe::/64 via ::, eth1, 00:23:16
C      fe80::/64 via ::, eth1, 01:25:04
K      ff00::/8 [0/256] via ::, eth1, 01:25:05



rtr1#show ipv6 ospf database grace

            OSPFv3 Router with ID (1.1.1.1) (Process 1)
```

R2 Before Graceful Restart

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:00:40    eth1       0


R2#sh ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:00:40    eth1       0
R2#
R2#show ipv6 ospf route
OSPFv3 Process (1)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  5ffe::/64                                         1
     directly connected, eth1, Area 0.0.0.0


R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:54:20
C      5ffe::/64 via ::, eth1, 00:12:16
```

```
C      fe80::/64 via ::, eth2, 01:17:21
K      ff00::/8 [0/256] via ::, eth1, 01:19:12


R2#show ipv6 ospf database grace

            OSPFv3 Router with ID (2.2.2.2) (Process 1)
```

R1 During graceful restart

```
R1#restart ipv6 ospf graceful grace-period 1000
% Warning : OSPF6D process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
R1#
R1#show running-config ospfv3
!
!
R1#
R1#show ipv6 ospf neighbor
```

R2 During graceful restart

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:02:25*    eth1       0


R2#show ipv6 ospf route
OSPFv3 Process (1)
Codes: C - connected, D - Discard, 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

   Destination                                    Metric
     Next-hop
C  5ffe::/64                                         1
     directly connected, eth1, Area 0.0.0.0


R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:57:26
C      5ffe::/64 via ::, eth1, 00:15:22
C      fe80::/64 via ::, eth2, 01:20:27
K      ff00::/8 [0/256] via ::, eth1, 01:22:18
```

```
R2#show ipv6 ospf database grace

            OSPFv3 Router with ID (2.2.2.2) (Process 1)

                Grace-LSA (Interface eth1)

  LS age: 65
  LS Type: Grace LSA
  Link State ID: 0.0.0.3
  Advertising Router: 1.1.1.1
  LS Seq Number: 0x80000001
  Checksum: 0x1660
  Length: 36

      Grace Period: 200
      Restart Reason:
           Software Restart
```

R1 After graceful restart

```
root@R1:/home/ocnos# cd /usr/local/sbin/
root@R1:/usr/local/sbin# ./ospf6d -d
bash-5.0$
bash-5.0$ exit
R1#
R1#show running-config ospfv3
!
router ipv6 ospf 1
 router-id 2.2.2.2
!
interface xe3
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
!
interface xe15
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
!
R1#
R1#show ipv6 ospf neighbor

Total number of full neighbors: 2
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
2.2.2.2           1   Full/DR     00:00:39    eth1       0
R1#
```

R2 After graceful restart

```
R2#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/Backup     00:00:34    eth1       0

R2#show ipv6 ospf route
OSPFv3 Process (1)
Codes: C - connected, D - Discard, 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
```

```
   Destination                                    Metric
     Next-hop
C  5ffe::/64                                          1
     directly connected, eth1, Area 0.0.0.0

R2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 02:07:14
C      5ffe::/64 via ::, eth1, 00:25:10
C      fe80::/64 via ::, eth2, 01:30:15
K      ff00::/8 [0/256] via ::, eth2, 01:30:17

R2#show ipv6 ospf database grace

       OSPFv3 Router with ID (2.2.2.2) (Process 1)
```

Open Shortest Path First Command Reference

CHAPTER 1 OSPFv2 Commands

This chapter provides an alphabetized reference for each of the OSPFv2 commands. It includes the following commands:

- area authentication
- area default-cost
- area interface
- area interface authentication
- area interface network-type
- area interface passive
- area filter-list
- area nssa
- area range
- area shortcut
- area stub
- area virtual-link
- auto-cost reference bandwidth
- bfd all-interfaces
- capability cspf
- capability lls
- capability opaque
- capability restart
- capability traffic-engineering
- capability vrf-lite
- clear ip ospf
- compatible rfc1583
- debug ospf
- debug ospf database-timer rate-limit
- debug ospf events
- debug ospf ifsm
- debug ip ospf graceful-restart
- debug ip ospf lfa
- debug ip ospf redist
- debug ip ospf retransmission
- debug ospf lsa
- debug ospf nfsm
- debug ospf nsm
- debug ospf packet
- debug ospf rib

© 2025 IP Infusion Inc. Proprietary

- debug ospf route
- default-information originate
- default-metric
- distance
- distribute-list
- enable db-summary-opt
- fast-reroute keep-all-paths
- fast-reroute tie-break
- host area
- ip ospf authentication
- ip ospf authentication-key
- ip ospf bfd
- ip ospf cost
- ip ospf database-filter
- ip ospf dead-interval
- ip ospf disable
- ip ospf fast-reroute per-prefix candidate disable
- ip ospf flood-reduction
- ip ospf hello-interval
- ip ospf multi-area
- ip ospf message-digest-key
- ip ospf mtu
- ip ospf mtu-ignore
- ip ospf network
- ip ospf priority
- ip ospf retransmit-interval
- ip ospf transmit-delay
- log-adjacency-changes
- max-concurrent-dd
- maximum-area
- max-metric
- neighbor
- network
- ospf abr-type
- ospf area-interface-config-mode
- ospf flood-reduction
- ospf restart helper
- ospf router-id
- overflow database

© 2025 IP Infusion Inc. Proprietary

- overflow database external
- passive-interface
- redistribute
- restart ip ospf graceful
- router ospf
- show debugging ospf
- show ip ospf
- show ip ospf border-routers
- show ip ospf database brief
- show ip ospf database detail
- show ip ospf igp-shortcut-lsp
- show ip ospf igp-shortcut-route
- show ip ospf interface
- show ip ospf multi-area-adjacencies
- show ip ospf neighbor
- show ip ospf route
- show ip ospf valid
- show ip ospf virtual-links
- show ip protocols
- show ip route fast-reroute
- shutdown
- snmp context-name
- snmp restart ospf
- summary-address
- timers lsa arrival
- timers spf exp
- timers throttle lsa

© 2025 IP Infusion Inc. Proprietary

area authentication

Use this command to enable authentication for an OSPF area. Specifying the area authentication sets the authentication to Type 1 authentication or simple text password authentication (details in RFC 2328). Setting up a Type 1 authentication configures a 64-bit field for that particular network. All packets sent on this network must have this configured value in their OSPF header. This allows only routers that have the same passwords to join the routing domain. Give all routers that are to communicate with each other through OSPF the same authentication password.

Use the ip ospf authentication-key command to specify a simple text password.

Use the ip ospf message-digest-key command to specify an MD5 password.

Use the `no` parameter to remove the authentication specification for an area.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) authentication
area (A.B.C.D|<0-4294967295>) authentication message-digest
no area (A.B.C.D|<0-4294967295>) authentication message-digest
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as 4-octet unsigned integer value.
`message-digest`	Enables MD5 authentication in the specified area ID.

Default

Null authentication

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 authentication message-digest

(config)#router ospf 100
(config-router)#no area 1 authentication
```

area default-cost

Use this command to specify a cost for the default summary route sent into a stub or NSSA area. This command provides the metric for the summary default route, generated by the area border router, into the NSSA or stub area. Use this option only on an area border router that is attached to the NSSA or stub area. Refer to the RFC 3101 for information on NSSA.

Use the `no` form of this command to remove the assigned default-route cost.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) default-cost <0-16777215>
no area (A.B.C.D|<0-4294967295>) default-cost
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`default-cost`	Indicates the cost for the default summary route used for a stub or NSSA area.
`<0-16777215>`	Stub's advertised default summary cost. The default is 1.

Default

By default, route cost is 1

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example sets the default-cost to 10 for area 1.

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 default-cost 10

(config)#router ospf 100
(config-router)#no area 1 default-cost
```

area interface

Use this command to enable OSPF routing on the current interface.

Use no form of this command to disable OSPF routing on the current interface.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface IFNAME
no area (A.B.C.D|<0-4294967295>) interface IFNAME
```

Parameters

A.B.C.D	OSPF area ID as an IPv4 address.
<0-4294967295>	OSPF area ID as a decimal value.
IFNAME	Specify the interface name.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1
```

area interface authentication

Use this command to send and receive OSPF packets with the specified authentication method on the current interface.

Use no form of this command to stop sending and receiving OSPF packets with the specified authentication method on the current interface.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface IFNAME authentication (null|message-
  digest|)
no area (A.B.C.D|<0-4294967295>) interface IFNAME authentication
```

Parameters

`A.B.C.D`	OSPF area ID as an IPv4 address.
`<0-4294967295>`	OSPF area ID as a decimal value.
`IFNAME`	Specify the interface name.
`null`	Use no authentication.
`message-digest`	Use message digest authentication.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 authentication null
```

area interface network-type

Use this command to set the OSPF network type on the current interface.

Use no form of this command to unset the OSPF network type on the current interface.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface IFNAME network-type (broadcast|non-
  broadcast|point-to-multipoint (non-broadcast|)|point-to-point)
no area (A.B.C.D|<0-4294967295>) interface IFNAME network-type
```

Parameters

`A.B.C.D`	OSPF area ID as an IPv4 address.
`<0-4294967295>`	OSPF area ID as a decimal value.
`IFNAME`	Specify the interface name.
`broadcast`	Sets the network type to broadcast.
`non-broadcast`	Sets the network type to NBMA.
`point-to-multipoint`	Sets the network type to point-to-multipoint..
`non-broadcast`	Sets the network type to NBMA.
`point-to-point`	Sets the network type to point-to-point.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 network-type point-to-point
```

area interface passive

Use this command to suppress sending hello packets on the current interface.

Use no form of this command to unsuppress sending Hello packets on the current interface.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface IFNAME passive
no area (A.B.C.D|<0-4294967295>) interface IFNAME passive
```

Parameters

A.B.C.D	OSPF area ID as an IPv4 address.
<0-4294967295>	OSPF area ID as a decimal value.
IFNAME	Specify the interface name.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced in OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 passive
```

area filter-list

Use this command to configure a filter to advertise summary routes on an Area Border Router (ABR).

This command suppresses incoming and outgoing summary routes between this area and other areas. You use this command in conjunction with the `prefix-list` and `access-list` commands.

Use the `no` form of this command to remove a filter.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) filter-list prefix WORD (in|out)
area (A.B.C.D|<0-4294967295>) filter-list access WORD (in|out)
no area (A.B.C.D|<0-4294967295>) filter-list prefix WORD
no area (A.B.C.D|<0-4294967295>) filter-list prefix WORD (in|out)
no area (A.B.C.D|<0-4294967295>) filter-list access WORD (in|out)
```

Parameters

`A.B.C.D`	OSPF area ID as an IPv4 address.
`<0-4294967295>`	OSPF area ID as a decimal value.
`prefix`	Use prefix list to filter summary.
`WORD`	Name of the prefix list.
`access`	Use access list to filter summary.
`WORD`	Name of the access list.
`in`	Filter routes from other areas into this area.
`out`	Filter routes from this area into other areas.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#access-list 1 deny 172.22.0.0
(config)#router ospf 100
(config-router)#area 1 filter-list access 1 in
```

area nssa

Use this command to set an area as a Not-So-Stubby-Area (NSSA). There are no external routes in an OSPF stub area, so you cannot redistribute from another protocol into a stub area. An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas. However, the external routes from other areas still do not enter the NSSA. You can configure an area to be a stub area or an NSSA, but not both.

This command simplifies administration when connecting a central site using OSPF to a remote site that is using a different routing protocol. You can extend OSPF to cover the remote connection by defining the area between the central router and the remote router as a NSSA.

Use the `no` form of this command to remove this designation.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) nssa {translator-role (candidate|always)|
  stabilityinterval < 0-2147483647>|no-redistribution|default-information-originate
  (metric < 0-16777214>|metric-type <1-2>|metric <0-16777214> metric-type <1-
  2>|metric-type< 1-2> metric <0-16777214>|)|no-summary}

area (A.B.C.D|<0-4294967295>) nssa {translator-role (candidate|always|never)|
  stabilityinterval < 0-2147483647>|no-redistribution|default-information-originate
  (metric < 0-16777214>|metric-type <1-2>|metric <0-16777214> metric-type <1-
  2>|metric-type< 1-2> metric <0-16777214>|)|no-summary}

no area (A.B.C.D|<0-4294967295>) nssa

no area (A.B.C.D|<0-4294967295>) nssa {translator-role|stability-interval|no-
  redistribution |default-information-originate (route-map |) |no-summary}
```

Parameters

`A.B.C.D` OSPF Area ID in IPv4 address format.

`<0-4294967295>` OSPF Area ID as a decimal value.

`translator-role`

NSSA-ABR translator role

- `candidate` Translate NSSA-LSA to Type-5 LSA if router is elected.
- `never` Do not translate NSSA-LSA to Type-5 LSA.
- `always` Always translate NSSA-LSA to Type-5 LSA.

`stability-interval`

Stability timer for a NSSA area. If an elected translator determines its services are no longer required, it continues to perform its duties for this time interval. This minimizes excess flushing of translated Type-7 LSAs and provides a more stable translator transition.

- `<0-2147483647>`

  Stability interval in seconds.

`no-redistribution`

Do not redistribute into the NSSA.

`default-information-originate`

Originate Type-7 default LSA into the NSSA.

`metric`	Specify metric for default routes.
`<0-16777214>`	Specify metric value.
`metric-type`	Specify metric type (see RFC 3101).
`<1-2>`	Specify metric type: 1: Type 1 external route 2: Type 2 external route
`route-map`	OSPF default Route map reference.
`WORD`	Pointer to route-map entries.
`no-summary`	Do not inject inter-area routes into the NSSA.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router ospf 100
(config-router)#area 3 nssa translator-role candidate no-redistribution
default-information-originate metric 34 metric-type 2
```

area range

Use this command to summarize OSPF routes at an area boundary. A single summary route is then advertised to other areas by the Area Border Routers (ABRs). Routing information is condensed at area boundaries and outside the area. If the network numbers in an area are assigned in a way such that they are contiguous, the ABRs can be configured to advertise a summary route that covers all the individual networks within the area that fall into the specified range.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) range (A.B.C.D/M | A.B.C.D A.B.C.D)
area (A.B.C.D|<0-4294967295>) range (A.B.C.D/M | A.B.C.D A.B.C.D) non-advertise
area (A.B.C.D|<0-4294967295>) range A.B.C.D A.B.C.D non-advertise
no area (A.B.C.D|<0-4294967295>) range (A.B.C.D/M | A.B.C.D A.B.C.D)
no area (A.B.C.D|<0-4294967295>) range (A.B.C.D/M | A.B.C.D A.B.C.D) not-advertise
```

Parameters

`A.B.C.D`	Area range prefix or length e.g. X.X.X.X/length
`A.B.C.D`	Area range prefix e.g. A.B.C.D
`<0-4294967295>`	OSPF Area ID as a decimal value.
`A.B.C.D/M`	The area range prefix and length.
`not-advertise`	Does not advertise this range.

Default

Range is advertised by default.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 range 192.16.0.0/24

(config)#router ospf 100
(config-router)#no area 1 range 192.16.0.0/24
```

area shortcut

Use this command to configure the short-cutting mode of an area. An area shortcut enables traffic to go through the non-backbone area with a lower metric whether or not an ABR router is attached to the backbone area.

Use the no form of this command to disable this function.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) shortcut (default|enable|disable)
no area (A.B.C.D|<0-4294967295>) shortcut
```

Parameters

A.B.C.D	OSPF Area ID in IPv4 address format.
<0-4294967295>	OSPF Area ID as a decimal value.
default	Sets default short-cutting behavior.
enable	Forces short-cutting through the area.
disable	Disables short-cutting through the area.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 shortcut default
(config)#router ospf 100
(config-router)#area 52 shortcut disable
(config)#router ospf 100
(config-router)#no area 42 shortcut
```

area stub

Use this command to define an area as a stub area. There are two stub area router configuration commands: the `stub` and `default-cost` commands. In all routers attached to the stub area, configure the area by using the `stub` option of the area command. For an area border router (ABR) attached to the stub area, use the `area default-cost` command.

Use the `no-summary` parameter with this command to define a totally stubby area. Define an area as a totally stubby area when routers in the area do not need to learn about summary LSAs from other areas.

Use the `no` form of this command to disable this function.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) stub
area (A.B.C.D|<0-4294967295>) stub no-summary
no area (A.B.C.D|<0-4294967295>) stub
no area (A.B.C.D|<0-4294967295>) stub no-summary
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`no-summary`	Stops an ABR from sending summary link advertisements into the stub area.

Default

By default, no stub area is defined.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 stub no-summary
```

area virtual-link

Use this command to configure a link between two backbone areas that are physically separated through other non-backbone area.

In OSPF, all non-backbone areas must be connected to a backbone area. If the connection to the backbone is lost, the virtual link repairs the connection. Configure virtual links between any two backbone routers that have an interface to a common non-backbone area. The protocol treats these routers joined by a virtual link as if they were connected by an unnumbered point-to-point network.

Configure the hello-interval to be the same for all routers attached to a common network. A short hello-interval results in the router detecting topological changes faster but also an increase in the routing traffic. The retransmit-interval is the expected round-trip delay between any two routers in a network. Set the value to be greater than the expected round-trip delay to avoid needless retransmissions.

The transmit-delay is the time taken to transmit a link state update packet on the interface. Before transmission, the link state advertisements in the update packet, are incremented by this amount. Set the transmit-delay to be greater than zero. Also, take into account the transmission and propagation delays for the interface. Include the transit area ID and the corresponding virtual link neighbor's router ID in each virtual link neighbor to properly configure a virtual link.

Use the `no` parameter with this command to remove a virtual link.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D authentication (message-
  digest|null|)
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D authentication-key LINE
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D message-digest-key <1-255> md5
  LINE
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {dead-interval <1-65535>|hello-
  interval <1-65535>|retransmit-interval <1-3600>|transmit-delay <1-3600>}
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {fall-over bfd}
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D authentication
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D authentication-key
no message-digest-key <1-255>
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {dead-interval |hello-
  interval |retransmit-interval |transmit-delay }
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D {fall-over bfd}
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`A.B.C.D`	Specify IP address of the virtual link neighbor.
`authentication`	Enable authentication on this virtual link
`message-digest`	Cryptographic authentication.

`null`	Null authentication.
`authentication-key`	
	Set authentication key.
`LINE`	Authentication key ID of 8 characters.
`message-digest-key`	
	Set message digest key.
`<1-255>`	Set message digest key.
`md5`	Specify the MD5 key.
`LINE`	MD5 key.
`dead-interval`	The interval during which no packets are received and after which the router acknowledges a neighboring router as off-line.
`<1-65535>`	The interval in seconds. The default is 40 seconds.
`hello-interval`	The interval the router waits before it sends a hello packet.
`<1-65535>`	The interval in seconds. The default is 10 seconds.
`retransmit-interval`	
	The interval the router waits before it retransmits a packet.
`<1-3600>`	The interval in seconds. The default is 5 seconds.
`transmit-delay`	The interval the router waits before it transmits a packet.
`<1-3600>`	The interval in seconds. The default is 1 second
`fall-over`	Specify fall-over detection.
`bfd`	Bidirectional Forwarding Detection (BFD)

Default

Default intervals:

Dead interval : 40 seconds

Hello interval: 10 seconds

Retransmit interval: 5 seconds

Transmit delay: 1 second

Command Mode

Router mode

OcNOS version 1.3

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 virtual-link 10.10.11.50 hello 5 dead 10
```

The below example shows the differences between with and without encrypted passwords.

0x - Already encrypted password

WORD - Authentication key (without encryption, ie doesn’t begin with 0x)

```
(config-router)#area 0.0.0.1 virtual-link 2.2.2.2 authentication
authentication-key 0x1234567891234567

(config-router)#area 0.0.0.1 virtual-link 2.2.2.2 authentication
authentication-key test
```

auto-cost reference bandwidth

Use this command to control how OSPF calculates the default metric for the interface.

Use the `no` form of this command to assign cost based only on the interface bandwidth.

Command Syntax

```
auto-cost reference-bandwidth <1-4294967>
no auto-cost reference-bandwidth
```

Parameters

`<1-4294967>`	The reference bandwidth in Mbps per second. The default is 100 Mbps.

Default

By default, OSPF calculates the OSPF metric for an interface by dividing the reference bandwidth by the interface bandwidth. The default value for the reference bandwidth is 100Mbps. The auto-cost command is used to differentiate high bandwidth links. For multiple links with high bandwidth, specify a larger reference bandwidth value to differentiate cost on those links.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#auto-cost reference-bandwidth 50

(config)#router ospf 100
(config-router)#no auto-cost reference-bandwidth
```

bfd all-interfaces

Use this command to enable Bidirectional Forwarding Detection (BFD) on all interfaces.

Use the `no` form of this command to disable BFD.

Command Syntax

```
bfd all-interfaces
no bfd all-interfaces
```

Parameters

None

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#bfd all-interfaces

(config)#router ospf 100
(config-router)#no bfd all-interfaces
```

capability cspf

Use this command to enable the CSPF (Constrained Shortest Path First) feature for an OSPFv2 or OSPFv3 instance.

Use the `no` parameter with this command to disable CSPF functionality for the OSPFv2 or OSPFv3 instance.

Command Syntax

```
capability cspf
no capability cspf
```

Parameters

None

Default

By default, CSPF functionality for the OSPFv2 or OSPFv3 instance is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#capability cspf

(config)#router ospf 100
(config-router)#no capability cspf
```

capability lls

Use this command to enable link-local signaling feature on OSPF router.

Use no parameter to disable link-local signaling feature on OSPF router.

Command Syntax

```
capability lls
no capability lls
```

Parameters

None

Default

By default, capability lls is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#capability lls

(config-router)#no capability lls
```

capability opaque

Use this command to enable opaque-LSAs which are Type 9, 10 and 11 LSAs that deliver information used by external applications.

Use the `no` form of this command to disable the feature.

Command Syntax

```
capability opaque
```

```
no capability opaque
```

Parameters

None

Default

By default, opaque-LSA is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#capability opaque

(config)#router ospf 100
(config-router)#no capability opaque
```

capability restart

Use this command to enable OSPF graceful restart or restart signaling. If a router is not restart-enabled, it cannot enter graceful restart mode and act as a helper.

Use the `no` parameter with this command to disable the features.

Command Syntax

```
capability restart graceful
no capability restart graceful
```

Parameters

`graceful`	Specify enabling OSPF graceful restart feature.

Default

By default, OSPF graceful restart or restart signaling is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#capability restart graceful

(config)#router ospf 100
(config-router)#no capability restart graceful
```

capability traffic-engineering

Use this command to enable traffic engineering feature on OSPF router.

Use `no` parameter to disable traffic engineering feature on OSPF router.

Command Syntax

```
capability traffic-engineering
no capability traffic-engineering
```

Parameters

None

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf
(config-router)#capability traffic-engineering

(config-router)#no capability traffic-engineering
```

capability vrf-lite

Use this command to apply multi-VRF capability to OSPF process or to decouple the PE router from the VPN backbone.

Use `no` parameter to deny multi-VRF capability to OSPF process or to avoid decoupling the PE router from the VPN backbone.

Command Syntax

```
capability vrf-lite
```

```
no capability vrf-lite
```

Parameters

None

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf
(config-router)#capability vrf-lite

(config-router)#no capability vrf-lite
```

clear ip ospf

Use this command to clear and restart all OSPF routing processes or a given OSPF routing process.

Command Syntax

```
clear ip ospf (<0-65535>|) process
```

Parameter

<0-65535>	Specify the process ID.

Command Mode

Privileged Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip ospf process
#clear ip ospf 555 process
```

compatible rfc1583

Use this command to restore the method used to calculate summary route costs per RFC.

RFC 1583 specified a method for calculating the metrics for summary routes based on the minimum metric of the component paths available. RFC 2328 specifies a method for calculating metrics based on maximum cost. With this change, it is possible that all of the ABRs in an area might not be upgraded to the new code at the same time. This command addresses this issue and allows the selective disabling of RFC 2328 compatibility.

Use the `no` parameter with this command to disable RFC 1583 compatibility.

Command Syntax

```
compatible rfc1583
```

```
no compatible rfc1583
```

Parameters

None

Default

By default, OSPF is RFC 2328 compatible

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#compatible rfc1583

(config)#router ospf 100
(config-router)#no compatible rfc1583
```

debug ospf

Use this command to specify debugging options for OSPF.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf (all|bfd|database-timer|events|ifsm|lsa|nfsm|nsm| packet|route|sr|)
debug ospf rib ({interface|redistribute}|)
no debug ospf (all|bfd|database-timer|events|ifsm|lsa|nfsm|nsm| packet|route|sr|)
no debug all ospf
no debug all
no debug ospf rib ({interface|redistribute}|)
```

Parameters

`all`	Enable or disable debugging for `ifsm`, `nsfm`, `lsa`, `nsm`, `events`, and `route`.
`bfd`	Debug Bidirectional Forwarding Detection (BFD)
`database-timer`	Debug OSPF rate-limiting values for LSA throttling (see debug ospf database-timer rate-limit)
`events`	Debug OSPF events information (see debug ospf events)
`ifsm`	Debug OSPF Interface State Machine (see debug ospf ifsm)
`lsa`	Debug OSPF Link State Advertisement (see debug ospf lsa)
`nfsm`	Debug OSPF Neighbor State Machine (see debug ospf nfsm)
`nsm`	Debug OSPF NSM information (see debug ospf nsm)
`packet`	Debug OSPF packets (see debug ospf packet)
`route`	Debug OSPF route information (see debug ospf route)
`rib`	Debug OSPF RIB information
`sr`	Debug OSPF segment routing information
`interface`	Debug OSPF RIB interface
`redistribute`	Debug OSPF RIB redistribute

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ospf all

#debug ospf bfd
#no debug ospf bfd
```

debug ospf database-timer rate-limit

Use this command to log when link-state advertisement (LSA) rate-limiting timers will expire. These messages are logged only when debug ospf lsa `generate` or debug ospf lsa `refresh` is enabled

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf database-timer rate-limit
no debug ospf database-timer rate-limit
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ospf database-timer rate-limit
```

debug ospf events

Use this command to specify debugging options for OSPF event troubleshooting. Use this command without parameters to turn on all the options.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf events ({abr|asbr|lsa|nssa|os|router|vlink}|)
no debug ospf events ({abr|asbr|lsa|nssa|os|router|vlink}|)
```

Parameters

`abr`	Debug OSPF ABR events.
`asbr`	Debug ASBR events.
`lsa`	Debug LSA events.
`nssa`	Debug NSSA events.
`os`	Debug OS interaction events.
`router`	Debug other router events.
`vlink`	Debug virtual link events.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ospf event abr
#debug ospf event asbr
#debug ospf event lsa
#no debug ospf event nssa
#debug ospf event os
#debug ospf event router
#debug ospf event vl
```

debug ospf ifsm

Use this command to specify debugging options for OSPF Interface Finite State Machine (IFSM) troubleshooting.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf ifsm ({events|status|timers}|)
no debug ospf ifsm ({events|status|timers}|)
```

Parameters

`events`	Debug IFSM event information
`status`	Debug IFSM status information
`timers`	Debug IFSM timer information

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ospf ifsm events
#debug ospf ifsm status
#debug ospf ifsm timers
```

debug ip ospf graceful-restart

Use this command to specify debugging option for OSPF graceful restart.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ip ospf graceful-restart (detail|terse|)
no debug ip ospf graceful-restart (detail|terse|)
```

Parameters

`detail`	Debug OSPF graceful restart detail information
`terse`	Debug OSPF graceful restart summary information

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf graceful-restart detail
```

debug ip ospf lfa

Use this command to specify the debugging options for OSPFv2 Loop-free Alternate path

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ip ospf lfa
no debug ip ospf lfa
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf lfa
```

debug ip ospf redist

Use this command to display debugging option for OSPF redistribute information

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ip ospf redist (detail|terse|)
no debug ip ospf redist (detail|terse|)
```

Parameters

`detail`	Debug OSPF redistribute detail information
`terse`	Debug OSPF redistribute summary information

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf redistribute detail
```

debug ip ospf retransmission

Use this command to display debug logs of OSPF retransmission information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ip ospf retransmission
no debug ip ospf retransmission
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf retransmission
```

debug ospf lsa

Use this command to specify debugging options for OSPF Link State Advertisements (LSA) troubleshooting.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf lsa ({flooding|generate|install|maxage|refresh}|)
no debug ospf lsa ({flooding|generate|install|maxage|refresh}|)
```

Parameters

`flooding`	Debug LSA flooding.
`generate`	Debug LSA generation.
`install`	Debug LSA installation.
`maxage`	Debug the maximum age processing.
`refresh`	Debug LSA refresh.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ospf lsa refresh
#debug ospf lsa flooding
#debug ospf lsa install
#debug ospf lsa maxage
#debug ospf lsa generate
```

debug ospf nfsm

Use this command to specify debugging options for OSPF Neighbor Finite State Machines (NFSMs).

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf nfsm ({events|status|timers}|)
no debug ospf nfsm ({events|status|timers}|)
```

Parameters

`events`	Debug NFSM event information
`status`	Debug NFSM status information
`timers`	Debug NFSM timer information

Command Mode

Privileged Exec mode Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ospf nfsm events
#no debug ospf nfsm timers
```

debug ospf nsm

Use this command to specify debugging options for OSPF NSM information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf nsm ({interface|redistribute}|)
no debug ospf nsm ({interface|redistribute}|)
```

Parameters

`interface`	Debug NSM interface information.
`redistribute`	Debug NSM redistribute information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The `debug ospf nsm` command enables the display of debug information related to NSM.

```
#debug ospf nsm interface
#no debug ospf nsm redistribute
```

debug ospf packet

Use this command to specify debugging options for OSPF packets.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf packet ({hello|dd|ls-request|ls-update|ls-ack|send|recv|detail}|)
no debug ospf packet ({hello|dd|ls-request|ls-update|ls-ack|send|recv|detail}|)
```

Parameters

`hello`	Debug OSPF hello packets.
`dd`	Debug OSPF database.
`ls-request`	Debug OSPF link state requests.
`ls-update`	Debug OSPF link state updates.
`ls-ack`	Debug OSPF link state acknowledgments.
`send`	Debug OSPF sent packets.
`recv`	Debug OSPF received packets.
`detail`	Debug OSPF detailed information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ospf packet detail
#debug ospf packet dd send detail
#no debug ospf packet ls-request recv detail
```

debug ospf rib

Use this command to display debug information about the interaction between the OSPF process and the Routing Information Base (RIB).

Use `no` parameter of this command to disable debugging output.

Command Syntax

```
debug ospf rib ({interface|redistribute}|)
no debug ospf rib ({interface|redistribute}|)
debug ip ospf redistribute
no debug ip ospf redistribute
```

Parameters

`interface`	Debug RIB interface information.
`redistribute`	Debug RIB redistribute information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#debug ospf rib interface
#no debug ospf rib redistribute
```

debug ospf route

Use this command to debug route calculation. Use this command without parameters to turn on all the options.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ospf route ({ase|ia|install|spf}|)
no debug ospf route ({ase|ia|install|spf}|)
```

Parameters

`ase`	Debug OSPF external route calculation.
`ia`	Debug OSPF Inter-Area route calculation.
`install`	Debug OSPF route installation.
`spf`	Debug OSPF SPF calculation.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ospf route
#no debug ospf route ia
#debug ospf route install
```

default-information originate

Use this command to create a default external route into an OSPF routing domain.

Use the `no` parameter with this command to disable this feature.

The system acts like an Autonomous System Boundary Router (ASBR) when you use the `default-information originate` command to redistribute routes into an OSPF routing domain. An ASBR does not by default generate a default route into the OSPF routing domain.

When you give the `default-information originate` command, also specify a `route-map` to avoid a dependency on the default network in the routing table.

Command Syntax

```
default-information originate
default-information originate {metric <0-16777214>|metric-type (1|2)|?route-map
  WORD|always}
no default-information originate
no default-information originate {metric|metric-type|?route-map|always}
```

Parameters

`always`	Used to advertise the default route regardless of whether there is a default route.
`metric`	Sets the OSPF metric used in creating the default route.
`<0-16777214>`	Sets the OSPF metric used in creating the default route. The default metric value is 10. The value used is specific to the protocol.
`metric-type`	The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101).
`1`	Sets OSPF External Type 1 metric.
`2`	Sets OSPF External Type 2 metric (default).
`route-map`	Route map.
`WORD`	Specify the name of route map.

Default

Sets the OSPF metric used in creating the default route. The default metric value is 10.

The value used is specific to the protocol. metric-type The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101).

By default, 2 Sets OSPF External Type 2 metric.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#default-information originate always metric 23 metric-type 2
route-map myinfo

(config)#router ospf 100
(config-router)#no default-information originate metric metric-type route-map
```

default-metric

Use this command to set a default metric for OSPF.

A default metric facilitates redistributing routes with incompatible metrics. If the metrics do not convert, the default metric provides an alternative. Use this command to use the same metric value for all redistributed routes. Use this command in conjunction with the redistribute command.

Use the `no` parameter with this command to return to the default state.

Command Syntax

```
default-metric <1-16777214>
no default-metric
```

Parameters

`<1-16777214>`	Default metric value.

Default

Built-in, automatic metric translations, as appropriate for each routing protocol.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#default-metric 100
```

distance

Use this command to set OSPF administrative distances.

The administrative distance rates the trustworthiness of a routing information source. A higher distance value means a lower trust rating. For example, an administrative distance of 255 means that the routing information source cannot be trusted and should be ignored.

Use the `no` form of this command to restore the default value (110).

Command Syntax

```
distance <1-255>
distance <1-255> A.B.C.D/M (WORD|)
distance ospf {intra-area <1-255>|inter-area <1-255>|external <1-255>}
no distance <1-255>
no distance <1-255> A.B.C.D/M (WORD|)
no distance ospf {intra-area |inter-area |external}
```

Parameters

`<1-255>`	Used alone, this parameter specifies a default administrative distance used when no other specification exists for a routing information source.
`intra-area`	Routes within an area.
`<1-255>`	Distance for all routes within an area
`inter-area`	Routes from one area to another area.
`<1-255>`	Distance for all routes from one area to another area.
`external`	Routes from other routing domains learned by redistribution.
`<1-255>`	Distance for routes from other routing domains learned by redistribution.
`A.B.C.D/M`	Distance for routes to prefixes whose nexthop matches this address.
`WORD`	Name of access list to apply to route updates.

Default

By default, distance for each type of route (intra-, inter-, or external) is 110

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#distance ospf inter-area 20 intra-area 10 external 40
```

distribute-list

Use this command to filter networks in routing updates. This command redistributes other routing protocols into the OSPF routing table.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
distribute-list WORD out (kernel|connected|static|rip|bgp|isis|ospf (<0-65535>|))
distribute-list WORD in
no distribute-list WORD out (kernel|connected|static|rip|bgp|isis|ospf (<0-
  65535>|))
no distribute-list WORD in
```

Parameters

`WORD`	Specify the name of the access list.
`in`	Filter incoming routing updates.
`out`	Filter outgoing routing updates.
`kernel`	Specify kernel routes.
`connected`	Specify connected routes.
`static`	Specify static routes.
`rip`	Specify RIP routes.
`bgp`	Specify BGP routes.
`isis`	Specify IS-IS routes.
`ospf`	Specify OSPF process.
`<0-65535>`	Specify OSPF process ID <1-65535>. If not specified, this command redistribute OSPF instance with process ID 0.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the distribution of BGP routing updates based on the access list `list1` (network 172.10.0.0).

```
#configure terminal
(config)#access-list list1 permit 172.10.0.0
(config)#router ospf 100
(config-router)#distribute-list list1 out bgp
```

```
(config-router)#redistribute bgp
```

enable db-summary-opt

Use this command to enable the database summary list optimization for OSPFv2.

When this feature is enabled, the database exchange process is optimized by removing the LSA from the database summary list for the neighbor, if the LSA instance in database summary list is the same as or less recent than the listed LSA in the database description packet received from the neighbor.

Use the `no` form of this command to disable database summary list optimization.

Command Syntax

```
enable db-summary-opt
```

```
no enable db-summary-opt
```

Parameters

None

Default

By default, database summary list optimization for OSPFv2 is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf
(config-router)#enable db-summary-opt
(config-router)#no enable db-summary-opt
```

fast-reroute keep-all-paths

Use this command to enable fast rerouting on all OSPF interfaces.

Use the `no` parameter with this command to disable fast rerouting.

Command Syntax

```
fast-reroute keep-all-paths
no fast-reroute keep-all-paths
```

Parameters

None

Defaults

By default, fast rerouting is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 200
(config-router)#fast-reroute keep-all-paths
```

fast-reroute tie-break

Use this command to set the tie-breaking policy for selecting a fast reroute repair path. You assign a priority to each type of repair path.

Use the `no` parameter with this command to set the tie-breaking policy for a specific type of repair path to its default priority. To set all types of repair paths to their default priorities, do not specify a repair path with the `no` form of this command.

Command Syntax

```
fast-reroute tie-break (primary-path|interface-disjoint|node-protecting|broadcast-
  interface-disjoint) index <1-10>
no fast-reroute tie-break
no fast-reroute tie-break (primary-path|interface-disjoint|node-
  protecting|broadcast-interface-disjoint) index <1-10>
```

Parameters

`primary-path` Use a path from the Equal-Cost Multipath Path (ECMP) set. An ECMP found during the primary shortest path first (SPF) repair might not be desirable in networks where traffic exceeds the capacity of any single link.

`interface-disjoint`

Do not select point-to-point interfaces that have no alternate next hop for rerouting if the primary gateway fails, thus protecting the interface.

`node-protecting`

Bypass the `primary-path` gateway router which might not protect the router that is the next hop in the primary path.

`broadcast-interface-disjoint`

Do not use the interface if connected to a broadcast network. Repair paths protect links when a repair path and a protected primary path use *different* next-hop interfaces. However, on broadcast interfaces, if the repair path is computed via the same interface as the primary path, but their next-hop gateways are different, the router is protected but the link might not be.

`index` Tie break priority. A lower value has higher preference.

`<1-10>` Range of priority values.

Defaults

By default, LFA backup path is calculated based on interface-disjoint.

The default priority scheme is:

1. primary-path
2. interface-disjoint
3. node-protecting
4. broadcast-interface-disjoint

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 200
(config-router)#fast-reroute tie-break interface-disjoint index 1
```

host area

Use this command to configure a stub host entry belonging to a particular area.

Using this command, you can advertise specific host routes in the router-LSA as stub link. Since stub host belongs to the specified router, specifying cost is not important.

Use the `no` form of this command to remove the host area configuration.

Command Syntax

```
host A.B.C.D area (A.B.C.D|<0-4294967295>)
host A.B.C.D area (A.B.C.D|<0-4294967295>) cost <0-65535>
no host A.B.C.D area (A.B.C.D|<0-4294967295>)
no host A.B.C.D area (A.B.C.D|<0-4294967295>) cost
```

Parameters

`A.B.C.D`	Specify IP address of the host.
`area`	Set the OSPF area ID
`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`cost`	Specify cost for stub host entry.
`<0-65535>`	Specify cost for stub host entry.

Default

No host entry is configured

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#host 172.16.10.100 area 1
(config-router)#host 172.16.10.101 area 2 cost 10
```

ip ospf authentication

Use this command to send and receive OSPF packets with the specified authentication method on the current interface.

Use the `no` parameter with this command to disable the authentication.

Command Syntax

```
ip ospf authentication (null|message-digest|)
ip ospf A.B.C.D authentication (null|message-digest|)
no ip ospf (A.B.C.D|) authentication
```

Parameters

`A.B.C.D`	The IP address of the interface.
`null`	Use no authentication.
`message-digest`	Use message digest authentication.

Default

No default value is specified

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In this example, interface `eth0` is configured to have no authentication. This will override any `text` or `MD5` authentication configured on this interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf authentication null
```

ip ospf authentication-key

Use this command to specify an OSPF authentication password for neighboring routers.

This command creates a password (key) that is inserted into the OSPF header when OcNOS originates packets. Assign a separate password to each network for different interfaces. All neighboring routers on the same network with the same password exchange OSPF routing data.

The key can be used only when authentication is enabled for an area with the `area authentication` command.

Simple password authentication allows a password to be configured for each area. Configure the routers in the same routing domain with the same password.

Use the `no` parameter with this command to remove an OSPF authentication password.

Command Syntax

```
ip ospf (A.B.C.D|) authentication-key WORD
no ip ospf (A.B.C.D|) authentication-key
```

Parameters

`A.B.C.D`	The IP address of the interface.
`authentication-key`	Specify the authentication password.
`WORD`	Specify the OSPF Encrypted password (key) of 8 bytes.

Default

By default, no password used when exchanging OSPF routing data

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, an authentication key `test` is created on interface `eth0` in area `0`. Note that first authentication is enabled for area `0`.

```
#configure terminal
(config)#router ospf 100
(config-router)#network 10.10.10.0/24 area 0
(config-router)#area 0 authentication
(config-router)#exit
(config)#interface eth0
(config-if)#ip ospf 3.3.3.3 authentication-key test

OcNOS#sh run int xe1
!
interface xe1
ip ospf authentication
ip ospf authentication-key 0x94eebee8c349a4b0
```

```
!
(config)#int xe1
(config-if)#ip ospf authentication-key 0x94eebee8c349a4b0 -
```
encrypted password of 16 characters which can be obtained from sh run <int>

ip ospf bfd

Use this command to enable Bidirectional Forwarding Detection (BFD).

Use this command with either the `no` or `disable` parameter to disable BFD.

Command Syntax

```
ip ospf bfd (disable|)
no ip ospf bfd (disable|)
```

Parameters

`disable`	Specify to disable BFD.

Default

By default, ip ospf bfd is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf bfd
```

ip ospf cost

Use this command to explicitly specify the cost of the link-state metric in a router-LSA.

The interface cost indicates the overhead required to send packets across an interface. This cost is stated in the Router-LSA's link. The cost is inversely proportional to the bandwidth of an interface. By default, the cost of an interface is calculated based on the bandwidth (10^8/ bandwidth). Use this command to set the cost manually.

Use the `no` parameter with this command to reset the cost to its default value.

Command Syntax

```
ip ospf (A.B.C.D|) cost <1-65535>
no ip ospf (A.B.C.D|) cost
```

Parameters

`A.B.C.D`	The IP address of the interface.
`<1-65535>`	Specify the link-state metric.

Default

By default, the cost of an interface is calculated based on the bandwidth (10^8/ bandwidth). The default cost value is 10.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the cost as 10 on interface `eth0` for IP address 10.10.10.50.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf 10.10.10.50 cost 10
```

ip ospf database-filter

Use this command to turn on the LSA database-filter for a particular interface.

OSPF floods new LSAs over all interfaces in an area, except the interface on which the LSA arrives. This redundancy ensures robust flooding. However, too much redundancy can waste bandwidth and might lead to excessive link and CPU usage in certain topologies, resulting in destabilizing the network. To avoid this, use this command to block flooding of LSAs over specified interfaces.

Use the `no` parameter with this command to turn off the filter.

Command Syntax

```
ip ospf (A.B.C.D|) database-filter all out
no ip ospf (A.B.C.D|) database-filter
```

Parameters

`A.B.C.D`	The IP address of the interface.

Default

Disabled, all outgoing LSAs are flooded to the interface.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf database-filter all out
```

ip ospf dead-interval

Use this command to set the interval during which the router waits to receive an OSPF hello packet from the neighbor before declaring the neighbor down. This value is advertised in the router's hello packets. It must be a multiple of hello-interval and be the same for all routers on a specific network.

Use the `no` parameter with this command to return to the default time. If you have configured this command specifying the IP address of the interface and want to remove the configuration, use the `no` parameter with the specified IP address (`no ip ospf dead-interval A.B.C.D`).

Command Syntax

```
ip ospf (A.B.C.D|) dead-interval <1-65535>
no ip ospf (A.B.C.D|) dead-interval
```

Parameters

`A.B.C.D`	The IP address of the interface.
`dead-interval`	Specify the interval.
`<1-65535>`	Specify the interval in seconds.

Default

By default, dead interval is 40 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows configuring dead-interval for 10 seconds on eth0 interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf dead-interval 10
```

ip ospf disable

Use this command to completely disable OSPF packet processing on an interface.

This command overrides the network command.

Use the `no` option with this command to return to the default setting.

Command Syntax

```
ip ospf disable all
no ip ospf disable all
```

Parameters

None

Default

By default, this feature is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf disable all
```

ip ospf fast-reroute per-prefix candidate disable

Use this command to prohibit the interface from being used as the next hop in a repair path.

Use the `no` option with this command to use the interface as the next hop in a repair path.

Command Syntax

```
ip ospf fast-reroute per-prefix candidate disable
no ip ospf fast-reroute per-prefix candidate disable
```

Parameters

None

Default

No default value is specified

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf fast-reroute per-prefix candidate disable
```

ip ospf flood-reduction

Use this command to enable flood reduction on an interface. When this command is configured, an LSA sent out on the interface is set with the DNA bit in the LSA age field. The LSA is not refreshed every refresh interval if there is no change in LSA. Only changed LSAs are sent out on the interface

Use the `no` option with this command to disable flood reduction on an interface.

Command Syntax

```
ip ospf flood-reduction
```

```
no ip ospf flood-reduction
```

Parameters

None

Default

By default, flood reduction on an interface is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ip ospf flood-reduction

(config-if)#no ip ospf flood-reduction
```

ip ospf hello-interval

Use this command to specify the interval between hello packets.

The hello-interval is advertised in the hello packets. Configure the same hello-interval for all routers on a specific network. A shorter hello interval ensures faster detection of topological changes but results in more routing traffic.

Use the `no` parameter with this command to return to the default time.

Command Syntax

```
ip ospf (A.B.C.D|) hello-interval <1-65535>
no ip ospf (A.B.C.D|) hello-interval
```

Parameters

`A.B.C.D`	The IP address of the interface.
`hello-interval`	Specify the interval.
`<1-65535>`	Specify the interval in seconds.

Default

By default, hello interval is 10 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the hello-interval for 3 seconds on interface eth0.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf hello-interval 3
```

ip ospf multi-area

Use this command to enable multi-area adjacency on point-to-point network and other network types. Multi-area adjacency establishes adjacency between the Area Border Routers (ABRs). The interface of the ABR where this command is configured, shall be associated with multiple areas.

Use the `no` parameter to disable multi-area adjacency on the given interface on point-to-point network.

Command Syntax

```
ip ospf <0-65535> multi-area (A.B.C.D|<0-4294967295>)(neighbor A.B.C.D |)
no ip ospf <0-65535> multi-area (A.B.C.D|<0-4294967295>)
no ip ospf <0-65535> multi-area (A.B.C.D|<0-4294967295>) neighbor
```

Parameters

`<0-65535>`	OSPF process ID.
`A.B.C.D`	OSPF area ID in IP address format.
`<0-4294967295>`	OSPF area ID as a decimal value.
`A.B.C.D`	Neighbor IP address.

Note: The neighbor parameter is mandatory to create redundant Multi-Area adjacencies in one Area.

Default

By default, multi-area adjacency is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth1
(config-if)#ip ospf 0 multi-area 1

(config-if)#no ip ospf 0 multi-area 1
```

ip ospf message-digest-key

Use this command to register an MD5 key for OSPF authentication.

Use the `no` parameter with this command to remove an MD5 key.

Message Digest Authentication is cryptographic authentication. A key (password) and key-id are configured on each router. The router uses an algorithm based on the OSPF packet, the key, and the key-id to generate a message digest that is appended to the packet.

Use this command for uninterrupted transitions between passwords. This is helpful for administrators who want to change the OSPF password without disrupting communication. The system begins a rollover process until all the neighbors have adopted the new password. This allows neighboring routers to continue communication while the network administrator is updating them with a new password. The router will stop sending duplicate packets once it detects that all of its neighbors have adopted the new password.

Maintain only one password per interface, removing the old password whenever you add a new one. This prevents the local system from continuing to communicate with the system that is using the old password. Removing the old password also reduces overhead during rollover. All neighboring routers on the same network must have the same password value to enable exchange of OSPF routing data.

Command Syntax

```
ip ospf (A.B.C.D|) message-digest-key <1-255> md5 WORD
no ip ospf (A.B.C.D|) message-digest-key <1-255>
```

Parameters

`A.B.C.D`	IPv4 address of the interface.
`message-digest-key`	
	Specify a key ID.
`<1-255>`	Specify a key ID.
`md5`	Specify a key (password).
`WORD`	Specify an encrypted password (key).

Default

By default, MD5 key for OSPF authentication is disabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows OSPF authentication on the interface eth0 when IP address has not been specified.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf authentication message-digest
(config-if)#ip ospf message-digest-key 1 md5 yourpass
```

```
(config-if)#sh run int xe1

!

interface xe1

 ip ospf authentication

 ip ospf authentication-key 0x94eebee8c349a4b0

 ip ospf message-digest-key 1 md5 0x94eebee8c349a4b0

!
```

```
(config-if)#ip ospf message-digest-key 1 md5 0x94eebee8c349a4b0
```
- encrypted password of 16 characters which can be obtained from sh run <int>

The following example shows OSPF authentication on the interface eth0 for the IP address 1.1.1.1. (If the interface has two IP addresses assigned-- 1.1.1.1 & 2.2.2.2, OSPF authentication will be enabled only for the IP address 1.1.1.1)

```
(config)#interface eth0
(config-if)#ip ospf 1.1.1.1 authentication message-digest
(config-if)#ip ospf 1.1.1.1 message-digest-key 2 md5 yourpass
```

ip ospf mtu

Use this command to set MTU size for OSPF to construct packets based on this value. Whenever OSPF constructs packets, it uses interface MTU size as Maximum IP packet size. This command forces OSPF to use the specified value overriding the actual interface MTU size.

This command does not configure the MTU settings in the kernel. OSPF does not recognize MTU size changes made in the kernel until the MTU size is updated through this command.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
ip ospf mtu <576-65535>
no ip ospf mtu
```

Parameters

`mtu`	Specify an MTU size.
`<576-65535>`	Specify an MTU size.

Default

By default, OSPF uses interface MTU derived from the kernel.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf mtu 1480
```

ip ospf mtu-ignore

Use this command to configure OSPF so that it does not check the MTU size during DD (Database Description) exchange.

Use the `no` form of this command to make OSPF check the MTU size during DD exchange.

Command syntax

```
ip ospf (A.B.C.D|) mtu-ignore
no ip ospf (A.B.C.D|) mtu-ignore
```

Parameters

`A.B.C.D`	IP address of the interface.

Default

By default, during the DD exchange process, OSPF checks the MTU size described in DD packets received from its neighbor. If the MTU size does not match the interface MTU, the neighbor adjacency is not established. Using this command makes OSPF ignore this check and allows establishing of adjacency regardless of MTU size in the DD packet.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-router)#ip ospf mtu-ignore
```

ip ospf network

Use this command to set the OSPF network type.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
ip ospf network (broadcast|non-broadcast|point-to-multipoint|point-to-point)
ip ospf network point-to-multipoint non-broadcast
no ip ospf network
```

Parameters

`broadcast`	Sets the network type to broadcast.
`non-broadcast`	Sets the network type to NBMA.
`point-to-multipoint`	Sets the network type to point-to-multipoint.
`non-broadcast`	Sets the network type to NBMA.
`point-to-point`	Sets the network type to point-to-point.

Default

By default, OSPF network type is broadcast

Command Mode

interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the network to `point-to-point` type on the `eth0` interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf network point-to-point
```

ip ospf priority

Use this command to set the router priority to determine the designated router (DR) for the network.

A router with the higher router priority becomes the DR. If the priority is the same for two routers, the router with the higher router ID takes precedence.

Only routers with a nonzero priority value are eligible to become the designated or backup designated router. Configure router priority for broadcast or NBMA networks only and not for point-to-point networks.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
ip ospf (A.B.C.D|) priority <0-255>
no ip ospf (A.B.C.D|) priority
```

Parameters

`A.B.C.D`	The IP address of the interface.
`priority`	Specify the router priority of the interface.
`<0-255>`	Specify the router priority of the interface. The default value is 1.

Default

By default, ip ospf priority is 1

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the OSPF priority value to 3 on the `eth0` interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf priority 3
```

ip ospf retransmit-interval

Use this command to specify the time between link-state advertisement (LSA) retransmissions for adjacencies belonging to the interface.

After sending an LSA to a neighbor, the router keeps the LSA until it receives an acknowledgement. If the router does not receive an acknowledgement during the retransmit interval, it retransmits the LSA. Set the retransmission interval value conservatively to avoid needless retransmission. The interval should be greater than the expected round-trip delay between two routers.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
ip ospf (A.B.C.D|) retransmit-interval <1-3600>
no ip ospf (A.B.C.D|) retransmit-interval
```

Parameters

`A.B.C.D`	The IPv4 address of the interface.
`retransmit-interval`	Specify the interval.
`<1-3600>`	Specify the interval in seconds.

Default

By default, retransmit interval is 5 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the `ospf retransmit interval` to 6 seconds on the eth0 interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf retransmit-interval 6
```

ip ospf transmit-delay

Use this command to set the estimated time it takes to transmit a link-state-update packet on the interface.

The transmit delay value adds a specified time to the age field of an update. If the delay is not added, the time in which the LSA transmits over the link is not considered. This command is especially useful for low speed links. Add transmission and propagation delays when setting the transmit delay value.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
ip ospf (A.B.C.D|) transmit-delay <1-65535>
no ip ospf (A.B.C.D|) transmit-delay
```

Parameters

`A.B.C.D`	The IPv4 address of the interface.
`transmit-delay`	Specify the time to transmit a link-state update.
`<1-65535>`	Specify the time in seconds to transmit a link-state update.

Default

By default, transmit delay is 1 second

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the OSPF transmit delay time to 3 seconds on the `eth0` interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ip ospf transmit-delay 3
```

log-adjacency-changes

Use this command for the router to send a SYSLOG message when an OSPF neighbor goes up or down.

Use `no` parameter of this command to stop sending SYSLOG message.

Command Syntax

```
log-adjacency-changes (brief|detail|)
no log-adjacency-changes
```

Parameters

`brief`	Sends a SYSLOG message for each state change, not just when a neighbor goes up or down.
`detail`	Sends a SYSLOG message for each state change, not just when a neighbor goes up or down.

Default

Default option is brief.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#log-adjacency-changes
(config-router)#log-adjacency-changes detail

(config-router)#no log-adjacency-changes
(config-router)#no log-adjacency-changes detail
```

max-concurrent-dd

Use this command to limit the number of Database Descriptors (DD) that can be processed concurrently.

This command is useful when a router's performance is affected from simultaneously bringing up several OSPF adjacencies. This command limits the maximum number of DD exchanges that can occur concurrently per OSPF instance, thus allowing for all of the adjacencies to come up.

Use the `no` option with this command to remove the limit.

Command Syntax

```
max-concurrent-dd <1-65535>
no max-concurrent-dd
```

Parameters

`<1-65535>`	Specify the number of DD processes.

Default

By default, max concurrent dd value is 64

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example set the `max-concurrent-dd` value to 4.

```
#configure terminal
(config)#router ospf 100
(config-router)#max-concurrent-dd 4
```

maximum-area

Use this command to configure the maximum number of OSPF areas.

Use the `no` parameter with this command to disable the limit.

Command Syntax

```
maximum-area <1-4294967294>
no maximum-area
```

Parameters

`<1-4294967294>` Specify the maximum number of OSPF areas.

Default

By default, ospf maximum area is 4294967294

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#maximum-area 5
```

max-metric

Use this command to advertise a maximum cost i.e 65535 for the LSA's that the router generates. This feature can be configured for in a router-LSA, summary-LSA, external-LSA and stub links.

It also provides support to configure an on-startup time. On-startup time is a one-time event which occurs when the node is up initially. The time configured will be the duration till which the router advertises its LSA’s with the maximum cost.

Based on the metric that is advertised, the neighboring OSPF routers re-trigger their SPF's, and calculate routes that don't pass through the new router.

Use the `no` parameter with this command to remove the configuration.

Command Syntax

```
max-metric router-lsa (on-startup (<5-86400>)|) ({ external-lsa (<1-16777215>|) |
  summary-lsa (<1-16777215>|) | include-stub }|)
no max-metric router-lsa (on-startup |) ({ external-lsa  | summary-lsa  | include-
  stub }|)
```

Parameters

```
router-lsa      router-lsa
external-lsa    External LSA max metric. <0-16777215>  External LSA max metric
                values
include-stub    Set the metric of a stub link in the router LSA to the default
                max-metric value
on-startup      Startup metric values for router LSA
summary-lsa     Summary LSA max metric. <0-16777215>  Summary LSA max metric
                values
```

Default

No default value.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 6.0.0.

Examples

```
OcNOS(config)#router ospf 1
OcNOS(config-router)#max-metric router-lsa summary-lsa 2000 external-lsa 300
include-stub
OcNOS(config-router)#commit
OcNOS(config-router)#
OcNOS(config-router)#do sh run ospf
!
router ospf 1
 max-metric router-lsa include-stub external-lsa 300 summary-lsa 2000
 redistribute connected
 network 100.0.0.0/24 area 0.0.0.0
```

```
 network 133.0.0.0/24 area 0.0.0.0
!

OcNOS(config-router)#no max-metric router-lsa
OcNOS(config-router)#commit
OcNOS(config-router)#do sh run ospf
!
router ospf 1
 redistribute connected
 network 100.0.0.0/24 area 0.0.0.0
 network 133.0.0.0/24 area 0.0.0.0
!

OcNOS(config-router)#max-metric router-lsa on-startup 300 external-lsa 200
OcNOS(config-router)#commit
OcNOS(config-router)#
OcNOS(config-router)#do sh run ospf
!
router ospf 1
 max-metric router-lsa
 max-metric router-lsa on-startup 300 external-lsa 200
 redistribute connected
 network 100.0.0.0/24 area 0.0.0.0
 network 133.0.0.0/24 area 0.0.0.0
!

OcNOS(config-router)#no max-metric router-lsa on-startup
OcNOS(config-router)#commit
OcNOS(config-router)#
OcNOS(config-router)#do sh run ospf
!
router ospf 1
 max-metric router-lsa
 redistribute connected
 network 100.0.0.0/24 area 0.0.0.0
 network 133.0.0.0/24 area 0.0.0.0
!
```

neighbor

Use this command to configure OSPF routers interconnecting to NBMA (Non-Broadcast Multi-Access) networks. Include one neighbor entry for each known non-broadcast network neighbor. Configure the neighbor address on the primary address of the interface.

Use the `no` parameter with this command to remove a configuration.

Command Syntax

```
neighbor A.B.C.D {cost <1-65535> | priority <0-255>| poll-interval <1-2147483647>}
no neighbor A.B.C.D { cost | priority | poll-interval}
no neighbor A.B.C.D
```

Parameters

`A.B.C.D`	Specify the interface IP address of the neighbor.
`priority`	Specify the router priority of the non-broadcast neighbor associated with the specified IP address.This parameter does not apply to point-to-multipoint interfaces.
`<0-255>`	Specify the router priority value of the non-broadcast neighbor associated with the specified IP address.
`poll-interval`	The reduced rate at which routers continue to send hello packets when a neighboring router has become inactive.
`<1-2147483647>`	Dead neighbor polling interval in seconds. Set this value much larger than hello interval.

Default

The default priority is 0 and polling interval is 120 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows neighbor configured with a priority value and poll interval time.

```
#configure terminal
(config)#router ospf 100
(config-router)#neighbor 1.2.3.4 priority 1 poll-interval 90
```

network

Use this command to enable OSPF routing with a specified area ID (and optionally an instance ID) on interfaces with IP addresses that match the specified network address.

OSPF routing is enabled per IPv4 subnet basis. You define the network address using the prefix length or a subnet mask.

Use the `no` parameter with this command to disable OSPF routing on the interfaces.

Command Syntax

Network address defined using the prefix length:

```
network A.B.C.D/M area (A.B.C.D|<0-4294967295>) (instance-id <0-255>|)
no network A.B.C.D/M area (A.B.C.D|<0-4294967295>) (instance-id)
```

Network address defined using subnet mask:

```
network A.B.C.D A.B.C.D area (A.B.C.D|<0-4294967295>) (instance-id <0-255>|)
no network A.B.C.D A.B.C.D area (A.B.C.D|<0-4294967295>) (instance-id)
```

Parameters

`A.B.C.D/M`	IPv4 network address with prefix length.
`A.B.C.D`	IPv4 network address.
`A.B.C.D`	Subnet mask where the bits on left side are set to 1 to represent the network part and the bits on the right side are set to 0 to represent the host part.
`area`	OSPF area ID
`A.B.C.D`	OSPF area ID in IPv4 address format.
`<0-4294967295>`	OSPF area ID as a decimal value.
`instance-id`	Instance ID
`<0-255>`	Instance ID value.

Default

No network area is configured

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following the use of the `network` command with OSPF multiple-instance support disabled.

```
#configure terminal
(config-router)#network 10.0.0.0/8 area 3
(config-router)#network 10.0.0.0/8 area 1.1.1.1
```

The following shows the use of the `network` command with OSPF multiple-instance support enabled.

```
(config)#router ospf 100
(config-router)#network 10.0.0.0/8 area 3 instance-id 4
```

ospf abr-type

Use this command to set an OSPF Area Border Router (ABR) type.

Use the `no` parameter with this command to revert the ABR type to the default setting (`cisco`).

Specifying the ABR type allows better functioning in a multi-vendor environment. The ABR types are:

- Cisco (RFC 3509): A router is considered an ABR if it has more than one area actively attached and one of them is the backbone area.
- IBM (RFC 3509): A router is considered an ABR if it has more than one area actively attached and the backbone area is configured. In this case the configured backbone need not be actively connected.
- Standard (RFC 2328): A router is considered an ABR if it has more than one area actively attached to it.

Command Syntax

```
ospf abr-type (cisco|ibm|standard|shortcut)
no ospf abr-type
```

Parameters

`cisco`	Specify an alternative ABR using Cisco implementation. This is the default ABR type.
`ibm`	Specify an alternative ABR using IBM implementation.
`standard`	Specify a standard ABR.
`shortcut`	Shortcut ABR.

Default

By default, ABR type is Cisco

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#ospf abr-type ibm
```

ospf area-interface-config-mode

Use this command to attach the interface to the router OSPF area. When this is enabled, any existing configuration of `router ospf` attached to the interface is erased.

Use no form of this command to disable. When this is disabled, any existing configuration of interface attached to the `router ospf` is erased.

Command Syntax

```
ospf area-interface-config-mode
```

```
no ospf area-interface-config-mode
```

Parameters

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS-OTN version 4.2.

Examples

```
#configure terminal
(config)#ospf area-interface-config-mode
```

ospf flood-reduction

Use this command to Enable flood reduction on all OSPF interface. When this command is configured, an LSA sent out on the OSPF interface is set with the DNA bit in the LSA age field. If there is no change in LSA, it is not refreshed every refresh interval. LSAs are sent out on the interface only if there is a change in an LSA

Use the `no` option with this command to disable flood reduction on all OSPF interfaces.

Command Syntax

```
ospf flood-reduction
```

```
no ospf flood-reduction
```

Parameters

None

Default

By default, flood reduction on all OSPF interfaces is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#ospf flood-reduction

(config-router)#no ospf flood-reduction
```

ospf restart grace-period

Use this command to set the grace period for restarting the router.

If graceful restart is enabled, NSM is notified about the grace period. If the OSPF daemon unexpectedly shuts down, NSM sends this value to the OSPF daemon when it comes up again which uses this value to end the graceful state.

Use the `no` parameter with this command to revert to the default.

Command Syntax

```
ospf restart grace-period <1-1800>
no ospf restart grace-period
```

Parameters

`grace-period`	Specify the grace period.
`<1-1800>`	Specify the grace period in seconds.

Default

The default grace period for restarting the OSPF router is 120 seconds.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ospf restart grace-period 250
```

ospf restart helper

Use this command to configure the helper behavior for graceful restart.

Use the `no` parameter with this command to revert to default.

Command Syntax

```
ospf restart helper max-grace-period <2-1800>
ospf restart helper never (router-id A.B.C.D|)
no ospf restart helper never
no ospf restart helper (never router-id (A.B.C.D | all) | max-grace-period|)
```

Parameters

`max-grace-period`	
	Help only if received grace-period is less than this value.
`<2-1800>`	
	Help only if received grace-period is less than this value.
`never`	Prevent the neighbor from entering helper mode.
`router-id`	Neighbor to never to act as helper.
`A.B.C.D`	Router ID of neighbor to never to act as helper.
`all`	All neighbors to never to act as helper.

Default

By default, router behave as helper. To disable it as helper, `ospf restart helper never` command should be configured. `ospf restart helper max-grace-period` – Max-grace-period to function as helper. If not configured, value will be the grace-period in restarting node.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ospf restart helper never router-id 1.1.1.1

#configure terminal
(config)#ospf restart helper only-reload max-grace-period 200

#configure terminal
(config)#no ospf restart helper never router-id all
```

ospf router-id

Use this command to specify a router ID for the OSPF process.

Configure each router with a unique router ID. In an OSPF router process which has active neighbors, a new router ID is used at the next reload or when you start the OSPF manually.

Use the `no` parameter with this command to force OSPF to use the previous router ID.

Command Syntax

```
ospf router-id A.B.C.D
router-id A.B.C.D

no ospf router-id
no router-id (A.B.C.D|)
```

Parameters

`A.B.C.D`	Specify the router ID in IPv4 address format.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows a specified router ID 2.3.4.5.

```
#configure terminal
(config)#router ospf 100
(config-router)#ospf router-id 2.3.4.5
```

overflow database

Use this command to limit the maximum number of LSAs that can be supported by the OSPF instance.

Use the `no` parameter with this command to have an unlimited number of LSAs.

Command Syntax

```
overflow database <0-4294967294> (hard|soft|)
no overflow database
```

Parameters

`<0-4294967294>`	The maximum number of LSAs
`hard`	Shutdown occurs if the number of LSAs exceeds the specified value.
`soft`	Warning message appears if the number of LSAs exceeds the specified value.

Default

No default value is specified. unlimited number of LSAs.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the database overflow to 5 and shutting down in that event.

```
#configure terminal
(config)#router ospf 100
(config-router)#overflow database 5 hard
```

overflow database external

Use this command to limits the number of AS-external-LSAs a router can receive once it is in the wait state.

Use the `no` parameter with this command to revert to default.

Command Syntax

```
overflow database external <0-2147483647> <0-65535>
no overflow database external
```

Parameters

`<0-2147483647>`	The maximum number of LSAs. This value should be the same on all routers in the AS.
`<0-65535>`	The number of seconds the router waits before trying to exit the database overflow state. If this parameter is 0, the router exits the overflow state only after an explicit administrator command.

Default

The default OSPF exit overflow interval is 0 second and number of external LSDB limit is unlimited.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the maximum number of LSAs to 5 and the time to recover from overflow state to be 3.

```
#configure terminal
(config)#router ospf 100
(config-router)#overflow database external 5 3
```

© 2025 IP Infusion Inc. Proprietary

passive-interface

Use this command to suppress sending Hello packets on all interfaces or on a specified interface.

This command configures OSPF on simplex Ethernet interfaces. Since a simplex interface represents only one network segment between two devices, configure the transmitting interface as a passive interface. This ensures that OSPF does not send hello packets for the transmitting interface. Both the devices can see each other via the hello packet generated for the receiving interface.

Use the `no` form with this command to resume sending `hello` packets on all interfaces, or on a specified interface.

Note: The command `passive-interface` alone configures all interfaces as non-passive.

Command Syntax

```
passive-interface
passive-interface IFNAME (enable|)
passive-interface IFNAME A.B.C.D (enable|)
no passive-interface
no passive-interface IFNAME
no passive-interface IFNAME A.B.C.D
```

Parameters

`enable`	Enable passive interface
`IFNAME`	The name of the interface.
`A.B.C.D`	IPv4 address of the interface.

Default

The default OSPF exit overflow interval is 0 second and number of external LSDB limit is 100000.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router ospf 100
(config-router)#passive-interface eth0 enable
```

redistribute

This command redistributes routes from a routing protocol, static route, and kernel route into an OSPF routing table.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
redistribute (connected|static|rip|bgp|isis|ospf (<1-65535>|)) {metric <0-
  16777214>|metric-type (1|2)|?route-map WORD|tag <0-4294967295>}
no redistribute (connected|static|rip|bgp|isis|ospf (<1-65535>|)) metric|metric-
  type|?route-map|tag}
```

Parameters

`connected`	Specify connected routes.
`static`	Specify static routes.
`rip`	Specify RIP routes.
`bgp`	Specify BGP routes.
`isis`	Specify IS-IS routes.
`ospf`	Specify OSPF process.
`<1-65535>`	Specify an OSPF process ID to redistribute a particular OSPF instance into another OSPF instance. If not specified, this command redistribute OSPF instance with process ID 0.
`metric`	Specify the external metric.
`<0-16777214>`	Specify the external metric.
`metric-type`	Specify the external metric-type (see RFC 3101):
`1`	Set OSPF External Type 1 metrics.
`2`	Set OSPF External Type 2 metrics.
`route-map`	Specify a route map reference.
`WORD`	Specify name of the route-map.
`tag`	Tag value to use as a “match” value for controlling redistribution via route maps
`<0-4294967295>`	Specify the route tag.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router ospf 100
(config-router)#redistribute bgp metric 12
```

The following example shows redistributing OSPF instance 2 into OSPF instance 1.

```
#configure terminal
(config)#router ospf 1
(config-router)#redistribute ospf 2
```

The following example shows redistributing OSPF instance `2` into OSPF instance `1`, with an external metric of `10`, metric type `1`, a route-map named `rmp1`, and an external route tag of `3`.

```
#configure terminal
(config)#router ospf 1
(config-router)#redistribute ospf 2 metric 10 metric-type 1 route-map rmp1 tag
3
```

restart ip ospf graceful

Use this command to restart OSPF gracefully.

After this command is executed, the router immediately shuts down. NSM is notified that OSPF has shut down gracefully. NSM preserves routes installed by OSPF until the grace period expires.

Command Syntax

```
restart ip ospf graceful (grace-period <2-1800>|)
```

Parameters

`grace-period`	Specify a grace period.
`<2-1800>`	Specify a grace period in seconds.

Default

By default, restart ospf graceful is disabled. Default value of grace-period is 120 seconds.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#restart ip ospf graceful grace-period 200
```

router ospf

Use this command to enter router mode and to configure an OSPF routing process.

Specify the process ID to configure multiple instances of OSPF. When running a single instance of OSPF, you do not need to specify a process ID.

Use the `no` parameter with this command to terminate an OSPF routing process.

Note: The CSPF database is maintained globally within OSPF. Removing any OSPF instance from a multi-instance setup impacts all CSPF LSP sessions and disrupts all ingress sessions. As a result, it is possible to restore only those sessions that still have a valid path after removing the OSPF instance.

Command Syntax

```
router ospf
router ospf <1-65535>

no router ospf
no router ospf <1-65535>
```

Parameters

`<1-65535>`	Process ID; should be unique for each routing process.

Default

No routing process defined

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows the use of the `router ospf` command to enter router mode. Note the change in the prompt.

```
#configure terminal
(config)#router ospf 100
(config-router)#
```

show debugging ospf

Use this command to display the set OSPF debugging option.

Command Syntax

```
show debugging ospf
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output from the `show debugging ospf` command. Some lines in this output wrap around, they might not wrap around in the actual display.

```
#show debugging ospf
OSPF debugging status:
  OSPF packet Link State Update debugging is on
  OSPF all events debugging is on
```

show ip ospf

Use this command to display general information about all OSPF routing processes.

Command Syntax

```
show ip ospf (<0-65535>|)
```

Parameters

<0-65535>	The ID of the router process for which information will be displayed. If this parameter is specified, only the information for the specified routing process is displayed.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip ospf 1
 Routing Process "ospf 1" with ID 4.1.1.1
 Process uptime is 1 minute
 Process bound to VRF default
 Conforms to RFC2328, and RFC1583 Compatibility flag is disabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 Supports Graceful Restart
 This router is an ASBR (injecting external routing information)
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 Refresh timer 10 secs
 Number of incomming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Initial LSA throttle delay 10 secs 0 msecs
 Minimum hold time for LSA throttle 20 secs 0 msecs
 Maximum wait time for LSA throttle 45 secs 0 msecs
 Minimum LSA arrival 1 secs 0 msecs
 Number of external LSA 5. Checksum 0x010632
 Number of opaque AS LSA 0. Checksum 0x000000
 Number of non-default external LSA 5
 External LSA database is unlimited.
 Number of LSA originated 6
 Number of LSA received 0
 Number of areas attached to this router: 1
    Area 0 (BACKBONE)
        Number of interfaces in this area is 1(1)
        Number of fully adjacent neighbors in this area is 0
        Area has no authentication
        SPF algorithm last executed 00:00:47.558 ago
```

```
SPF algorithm executed 2 times
Number of LSA 1. Checksum 0x0041e0
```

OSPF Routing Process Fields

Table 1-13 explains the routing process fields.

Table 1-13: show ip ospf output details

Field	Description
Routing Process with ID	OSPF process identifier and router identifier.
Process is not up	OSPF process is not running.
Process uptime is	OSPF process running time.
Process bound to VRF	VRF name
Router is in Graceful Restart	When in graceful restart.
Router is in Restart Signaling	When in restart signalling.
Bidirectional Forwarding Detection is configured	When BFD is enabled.
Conforms to RFC2328, and RFC1583 Compatibility flag is enabled	RFC compatibility.
Supports only single TOS (TOS0) routes	OSPF TOS-based routing was never deployed.
Supports opaque LSA	When opaque LSAs are supported.
Do not support Restarting	When restart is not supported.
Supports Graceful Restart	Method of restart in process.
Supports Restart Signaling	Method of restart in signaling process.
This router is an ABR, ABR Type is	Type of ABR: Standard (RFC2328) Alternative Cisco (RFC3509) Alternative IBM (RFC3509) Alternative Shortcut
This router is an ASBR (injecting external routing information)	Type of router function in the process.
SPF schedule delay initial	Initial SPF schedule delay.
SPF schedule delay min	Minimum delay between receiving a change to SPF calculation.
SPF schedule delay max	Maximum delay between receiving a change to SPF calculation.
Refresh timer	LSA refresh interval.
Number of incoming current DD exchange neighbors	Incoming neighbor Database Descriptors and maximum concurrent DDs.
Number of outgoing current DD exchange neighbors	Outgoing neighbor Database Descriptors and maximum concurrent DDs.

Table 1-13: show ip ospf output details (Continued)

Field	Description
Initial LSA throttle delay	Initial delay for the generation of LSAs.
Minimum hold time for LSA throttle	Minimum hold time between generation of the same LSA.
Maximum wait time for LSA throttle	Maximum wait time between generation of the same LSA.
Minimum LSA arrival	Minimum time between reception of new LSAs during flooding.
Number of external LSA	Number of AS external LSAs and checksum.
Number of opaque AS LSA	Number of AS opaque LSAs and checksum.
Number of non-default external LSA	For database overflow, number of non-default external LSAs.
External LSA database is unlimited	When the external LSA database is unlimited.
External LSA database limit	Maximum number of LSAs in database.
Exit database overflow state interval is	Exit database overflow state interval.
Exit database overflow state interval is not configured	When the exit database overflow state interval is not set.
OSPF is [not] in database overflow state now	Whether OSPF is in database overflow state now.
Next attempt to exit database overflow state in	How long until OSPF tries to exit the database overflow state.
LSDB database overflow limit	Maximum number of LSAs that can be supported by the OSPF instance.
LSDB exceed overflow limit	Whether OSPF is exceeding the maximum number of LSAs.
Number of LSA originated	LSAs originated by the OSPF instance.
Number of LSA received	LSAs received by the OSPF instance.
Number of areas attached to this router	As stated
Next fields are repeated for each area	As stated
Area	Area identifier.
(BACKBONE)	Area is a backbone.
no-summary	Area is a stub and does no import summaries.
(Inactive)	Area is not active.
Number of interfaces in this area is	Number of interfaces in this area.
Number of fully adjacent neighbors in this area	As stated.
Number of fully adjacent virtual neighbors through this area	As stated.
Area has no authentication	Area does not use authentication.
Area has simple password authentication	Area uses password authentication.

Table 1-13: show ip ospf output details (Continued)

Field	Description
Area has message digest authentication	Area uses MD5 authentication.
SPF algorithm last executed	As stated.
SPF algorithm executed	As stated.
Number of LSA	Number of LSAs in area link-state database and checksum.
End of area field	As stated.
NSSA Translator Role is	candidate: Translate Type-7 LSAs to Type-5 if router is elected. never: Do not translate Type-7 LSAs to Type-5. always: Always translate Type-7 LSAs to Type-5.
NSSA Translator State is	disabled: Router is not a border router. enabled: Router is a border router. elected: Router has been elected to be an NSSA translator.
Stability Interval	If an elected translator determines its services are no longer required, how long it continues to perform its services.
Number of NSSA Translator Events	As stated.
Shortcutting mode	Shortcut ABR that installs inter-area routes through non-backbone areas if non-backbone paths are really better: Default Enabled Disabled
S-bit consensus	Whether other ABR agrees on S-bit: ok no
Dste Status	Whether DSTE is enabled or disabled.

show ip ospf border-routers

Use this command to display the ABRs and ASBRs for OSPF instances.

Command Syntax

```
show ip ospf (<0-65535>|) border-routers
```

Parameters

`<0-65535>`	The ID of the router process for which information will be displayed.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output from the `show ip ospf border-routers` command.

```
#show ip ospf border-routers
OSPF process 1 internal Routing Table
Codes: i - Intra-area route, I - Inter-area route
i 10.15.0.1 [10] via 10.10.0.1, eth0, ASBR, Area 0.0.0.0
i 172.16.10.1 [10] via 10.10.11.50, eth1, ABR, ASBR, Area 0.0.0.0
```

Border Router Fields

Table 1-14 explains the border router fields.

Table 1-14: border router output details

Field	Description
Code	i: Intra-area route I: Inter-area route
Router ID	Router identifier of the destination
Cost	Cost of using this route.
via	Next hop IP address toward the destination.
is directly connected	Destination is directly connected.
Interface	Outgoing interface name.
Type	Router type of the destination: ABR or ASBR.
through Transit Area	Next hop is an area that carries traffic that neither originates nor terminates in the area itself.
Area	Area identifier from which this route was learned.
Transit	Area is a transit area.

show ip ospf database brief

Use this command to display a summary of the OSPF database.

Command Syntax

```
show ip ospf database (self-originate|max-age|adv-router A.B.C.D|)
show ip ospf <0-65535> database(self-originate|max-age|adv-router A.B.C.D|)
```

Parameters

`self-originate`	Self-originated link states.
`max-age`	LSAs which have reached the maximum age (3600 seconds).
`A.B.C.D`	IPv4 address of the advertising router.
`<0-65535>`	ID of the router process

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip ospf database
       OSPF Router process 100 with ID (100.100.100.72)
                Router Link States (Area 0.0.0.0)
Link ID         ADV Router      Age   Seq#       CkSum  Link count
10.100.12.57    10.100.12.57     930  0x80000003 0x90de 2
100.100.100.72  100.100.100.72   933  0x80000004 0x7592 2
                Net Link States (Area 0.0.0.0)
Link ID         ADV Router      Age   Seq#       CkSum
10.100.10.72    100.100.100.72   933  0x80000001 0x0bef
                Summary Link States (Area 0.0.0.0)
Link ID         ADV Router      Age   Seq#       CkSum  Route
10.60.0.0       10.100.12.57     928  0x80000001 0x5108 10.60.0.0/24
71.87.120.0     10.100.12.57     928  0x80000001 0xc2c5 71.87.120.0/24
127.0.0.1       10.100.12.57     928  0x80000001 0x23fb 127.0.0.1/32
```

OSPF Database Fields

Table 1-15 explains the fields for each database entry.

Table 1-15: ospf database output details

Field	Description
Link ID	The meaning of this field depends on the type of Link-State Advertisement (LSA). Type 1: Router LSA (depends on the type of network to which the router connects): Point-to-point network: neighbor's router ID. Transit network: IP address of the designated router's interface. Stub network: IP network or subnet address Virtual link: Neighbor's Router ID. Type 2: Network LSA: The IP address of the designated router's interface. Type 3: Summary LSA: The IP address of the network or subnet being advertised.
ADV Router	The ID of the router advertising the LSA.
Age	The age of the LSA.
Seq#	The sequence number of the LSA. This number increments each time a new instance of the LSA originates. This update helps other routers identify the most recent instance of the LSA.
CkSum	The fetch checksum of the complete LSA except the Age field.
Link count	Total number of links.
Route	Summary prefix address.

show ip ospf database detail

Use this command to display details of the OSPF database.

Command Syntax

```
show ip ospf database (asbr-summary|external|network|router|summary|nssa-
  external|opaque-link|opaque-area|opaque-as) (self-originate|adv-router A.B.C.D|)
show ip ospf <0-65535> database (asbr-summary|external|network|router|summary)
  (self-originate|adv-router A.B.C.D|)
show ip ospf database (asbr-summary|external|network|router|summary|nssa-
  external|opaque-link|opaque-area|opaque-as) A.B.C.D (self-originate|adv-router
  A.B.C.D|)
show ip ospf <0-65535> database (asbr-summary|external|network|router|summary|nssa-
  external|opaque-link|opaque-area|opaque-as) A.B.C.D (self-originate|adv-router
  A.B.C.D|)
```

Parameters

`<0-65535>`	The ID of the router process for which information should be displayed.
`asbr-summary`	Type 4 ASBR summary LSAs.
`external`	Type 5 external LSAs.
`network`	Type 2 network LSAs.
`router`	Type 1 router LSAs.
`summary`	Type 3 summary LSAs.
`nssa-external`	Type 7 NSSA external LSAs.
`opaque-link`	Type 9 LSAs which are not flooded beyond the local network.
`opaque-area`	Type 10 LSAs which are not flooded beyond the borders of their area.
`opaque-as`	Type 11 LSAs which are flooded throughout the AS.
`A.B.C.D`	Link state ID as an IP address.
`self-originate`	Display self-originated link states.
`adv-router`	Advertising router link states.
`A.B.C.D`	IPv4 address of advertising router.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example: external and self-originate Parameters

This is sample output with the `external` and `self-originate` parameters.

```
#show ip ospf database external self-originate

 OSPF Router process 100 with ID (10.10.11.50)
```

```
 AS External Link States
  LS age: 298
  Options: 0x2 (*|-|-|-|-|-|E|-)
  LS Type: AS-external-LSA
  Link State ID: 10.10.100.0 (External Network Number)
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0x7033
  Length: 36
  Network Mask: /24
 Metric Type: 2 (Larger than any link state path)
 TOS: 0
 Metric: 20
 Forward Address: 10.10.11.50
 External Route Tag: 0
```

Example: opaque-as and self-originate Parameters

This is sample output with the `opaque-as` and `self-originate` parameters.

```
#show ip ospf database opaque-as self-originate
 OSPF Router process 100 with ID (10.10.11.50)
 AS-Global Opaque-LSA
  LS age: 325
  Options: 0x2 (*|-|-|-|-|-|E|-)
  LS Type: AS-external Opaque-LSA
  Link State ID: 11.10.9.23 (AS-external Opaque-Type/ID)
  Opaque Type: 11
  Opaque ID: 657687
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0xb018
 Length: 25
```

Example: adv-router Parameter

This is a sample output with the `adv-router` parameter.

```
#show ip ospf database nssa-external adv-router 10.10.11.50
 OSPF Router process 100 with ID (10.10.11.50)
 NSSA-external Link States (Area 0.0.0.0)
 NSSA-external Link States (Area 0.0.0.1 [NSSA])
  LS age: 78
  Options: 0x0 (*|-|-|-|-|-|-|-)
  LS Type: AS-NSSA-LSA
  Link State ID: 0.0.0.0 (External Network Number For NSSA)
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0xc9b6
  Length: 36
  Network Mask: /0
 Metric Type: 2 (Larger than any link state path)
 TOS: 0
 Metric: 1
 NSSA: Forward Address: 0.0.0.0
OSPF Router process 100 with ID (10.10.11.50)
 NSSA-external Link States (Area 0.0.0.0)
 NSSA-external Link States (Area 0.0.0.1 [NSSA])
```

```
  LS age: 78
  Options: 0x0 (*|-|-|-|-|-|-|-)
  LS Type: AS-NSSA-LSA
 Link State ID: 0.0.0.0 (External Network Number For NSSA)
```

Example: router and Link State ID Parameters

This is sample output with the `router` and link state ID parameters.

```
#show ip ospf database router 10.10.11.50
 OSPF Router process 100 with ID (10.10.11.50)
 Router Link States (Area 0.0.0.0)
  LS age: 878
  Options: 0x2 (*|-|-|-|-|-|E|-)
  Flags: 0x3 : ABR ASBR
  LS Type: router-LSA
  Link State ID: 10.10.11.50
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000004
  Checksum: 0xe39e
  Length: 36
Number of Links: 1
 Link connected to: Stub Network
 (Link ID) Network/subnet number: 10.10.10.0
 (Link Data) Network Mask: 255.255.255.0
Number of TOS metrics: 0
 TOS 0 Metric: 10
 Router Link States (Area 0.0.0.1)
  LS age: 877
  Options: 0x2 (*|-|-|-|-|-|E|-)
  Flags: 0x3 : ABR ASBR
  LS Type: router-LSA
  Link State ID: 10.10.11.50
  Advertising Router: 10.10.11.50
 LS Seq Number: 80000003
```

Example: adv-router Parameter

This is sample output using the `adv-router` parameter for flood reduction.

```
#show ip ospf database summary adv-router 10.10.11.50

 OSPF Router process 100 with ID (10.10.11.50)
 Summary Link States (Area 0.0.0.0)
  LS age: 1(DoNotAge)
  Options: 0x2 (*|-|-|-|-|-|E|-)
  LS Type: summary-LSA
  Link State ID: 10.10.11.0 (summary Network Number)
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0x36ac
  Length: 28
  Network Mask: /24
 TOS: 0 Metric: 10
Summary Link States (Area 0.0.0.1)
  LS age: 989
  Options: 0x2 (*|-|-|-|-|-|E|-)
LS Type: summary-LSA
  Link State ID: 10.10.11.0 (summary Network Number)
```

```
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0x36ac
  Length: 28
  Network Mask: /24
 TOS: 0 Metric: 10

#show ip ospf database external self-originate

 OSPF Router process 100 with ID (10.10.11.50)

 AS External Link States
  LS age: 298
  Options: 0x2 (*|-|-|-|-|-|E|-)
  LS Type: AS-external-LSA
  Link State ID: 10.10.100.0 (External Network Number)
  Advertising Router: 10.10.11.50
  LS Seq Number: 80000001
  Checksum: 0x7033
  Length: 36
  Network Mask: /24
 Metric Type: 2 (Larger than any link state path)
 TOS: 0
 Metric: 20
 Forward Address: 10.10.11.50
 External Route Tag: 0
```

Database Detail Header Fields

Table 1-16 explains the fields for each database entry.

Table 1-16: ospf database detail header fields

Field	Description
LS age	Age of the LSA in seconds. “Do Not Age” is displayed if the DNA bit is set.
Options	LSA options as explained in Table 1-17.
Flags	ABR: Area border router ASBR: AS boundary router VL-endpoint: Endpoint of an active virtual link that is using the described area as a transit area Shortcut: shortcut ABR NSSA-Translator: NSSA border router with NSSA Translate or State enabled
LS Type	Type of LSA: Router-LSA Network-LSA Summary-LSA ASBR-summary-LSA AS-external-LSA AS-NSSA-LSA Link-Local Opaque-LSA Area-Local Opaque-LSA AS-external Opaque-LSA

Table 1-16: ospf database detail header fields

Field	Description
Link State ID	Identifier of the router described by the LSA.
Opaque Type	Opaque type used to identify the application type of the LSA: 9: link-local scope 10: area-local scope 11: LSA flooded throughout the AS
Opaque ID	Identifier used to differentiate LSAs of the same type.
Advertising Router	Identifier of the router that originated the LSA.
LS Seq Number	Sequence number of the LSA. This number increments each time a new instance of the LSA originates. This update helps other routers identify the most recent instance of the LSA.
Checksum	Checksum of the entire LSA, except the LS age field.
Length	Length of the LSA
I LSA	Indication LSA: ASBR set the infinity metric to tell all routers in the backbone not to originate DNA LSAs.

OSPF LSA Option Bits

Table 1-17 explains the fields for each database entry.

Table 1-17: ospf LSA option bits output details

Bit	Description
DN	Used in L3 VPNs. When a route learned from a customer network via OSPF is advertised across a BGP/VPN using Multiprotocol BGP and advertised back to a customer network via OSPF, a loop can happen where the OSPF route is redistributed back to the VPN service provider network via BGP. The DN-bit prevents this type of routing loop. When an OSPF router receives a Type 3, 5, or 7 LSA with the DN-bit set, it foes use that LSA for OSPF route calculations.
O	Originating router supports Type 9, 10, and 11 Opaque LSAs.
DC	Originating router supports OSPF over Demand Circuits.
L	Whether the OSPF packet contains a Link-Local Signaling (LLS) data block. This bit is set only in Hello and database description packets.
N/P	The N-bit is used only in Hello packets when the originating router supports Type-7 NSSA-External-LSAs. Neighboring routers with mismatched N-bit will not form a neighbor relationship. This restriction ensures that all OSPF routers within an area support NSSA capabilities. When the N-bit is set, the E-bit must be 0. The P-bit is used only in Type-7 NSSA-External-LSA headers. Due to this reason, the N- and P-bits share the same position in the options field. The P (Propagate) bit is set to inform an NSSA ABR to translate Type-7 LSAs into Type-5 LSAs.
MC	Originating router supports multicast extensions to OSPF (MOSPF)

Table 1-17: ospf LSA option bits output details (Continued)

Bit	Description
E	Originating router accepts AS External LSAs. The bit is set in all AS External LSAs and in all LSAs originated in the backbone and non-stub areas; and is be set to 0 in all Hellos and LSAs originated within a stub area. Additionally, this bit is used in Hello packets to indicate the capability of a router interface to send and receive Type-5 AS-External-LSAs. Neighboring routers with mismatched E-bit do not form a neighbor relationship. This restriction ensures that all OSPF routers within an area support the stub capabilities.
T	Originating router supports Multitopology OSPF (MT-OSPF.) Older OSPF specifications used this bit when the originating router support TOS-based routing. However, OSPF TOS-based routing was never deployed; therefore the T-bit was never used.

Type 1 Router LSAs (“router” Parameter)

Table 1-18 explains the fields for each database entry.

Table 1-18: router LSAs

Field	Description
Number of Links	Number of router links the LSA describes.
Link connected to	Description of the router link: another Router (point-to-point) a Transit Network Stub Network a Virtual Link
(Link ID)	Identifier of the router to which the link connects: Neighboring Router ID Designated Router address Network/subnet number Neighboring Router ID
(Link Data)	Extra information: Router Interface address Network Mask
Number of TOS metrics	Number of TOS (Type of Service) metrics for this link, not including the metric for TOS 0.
TOS 0 Metric	Cost of using this router link for TOS 0.

Type 2 Net Link States (“network” Parameter)

Table 1-19 explains the fields for each database entry.

Table 1-19: net LSAs

Field	Description
Network Mask	IP address mask for the network.
Attached Router	Identifiers of each router attached to the network.

Type 3 Summary LSAs (“summary” Parameter) and Type 4 ASBR Summary LSAs (“asbr-summary” Parameter)

Table 1-20 explains the fields for each database entry.

Table 1-20: summary and ASBR summary link states

Field	Description
Network Mask	For Type 3 LSAs, the destination network's IP address mask. Not meaningful for Type 4 link state advertisements.
TOS: 0 Metric	Cost of using this router link for TOS 0.

Type 5 AS External LSAs (“external” Parameter)

Table 1-21 explains the fields for each database entry.

Table 1-21: external LSAs

Field	Description
Network Mask	IP address mask for the advertised destination
Metric Type	1: Type 1 external metric that is comparable directly (without translation) to the link state metric 2: Type 2 external metric that is considered larger than any link state path
TOS	Always zero.
Metric	The cost of this route.
Forward Address	Data traffic for the advertised destination is forwarded to this address.
External Route Tag	Custom field attached to each external route whose use is defined by the application.

Type 7 NSSA External Link States (“nssa-external” Parameter)

Table 1-22 explains the fields for each database entry.

Table 1-22: NSSA external LSAs

Field	Description
Network Mask	IP address mask for the advertised destination
Metric Type	1: Type 1 external metric that is comparable directly (without translation) to the link state metric 2: Type 2 external metric that is considered larger than any link state path
Metric	The cost of this route.
NSSA: Forward Address	Data traffic for the advertised destination is forwarded to this address.
External Route Tag	Custom field attached to each external route whose use is defined by the application.

show ip ospf igp-shortcut-lsp

Use this command to show the IGP shortcut LSP used by OSPF.

Command Syntax

```
show ip ospf igp-shortcut-lsp
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip ospf igp-shortcut-lsp
 Tunnel-endpoint     Tunnel-id     Tunnel-metric
 8.8.8.8             101           2              active
```

Table 1-23 explains the fields in the output.

Table 1-23: show ip ospf igp-shortcut-lsp output details

Field	Description
Tunnel-endpoint	Tunnel endpoint address of ospf.
Tunnel-id	Tunnel address (destination port) for the session.
Tunnel-metric	Number of tunnel-metric.
active/inactive	Whether the tunnel is active or inactive.

show ip ospf igp-shortcut-route

Use this command to show the IGP shortcut route calculated by OSPF.

Command Syntax

```
show ip ospf (<0-65535>|) igp-shortcut-route
```

Parameters

<0-65535>	ID of the router process.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip ospf igp-shortcut-route
OSPF process 0:
8.8.8.8/32 [2] tunnel-id: 101, 8.8.8.8
15.15.15.15/32 [0] tunnel-id: 101, 8.8.8.8
20.20.15.0/24 [0] tunnel-id: 101, 8.8.8.8
```

Table 1-24 explains the fields in the output.

Table 1-24: show ip igp-shortcut-route output details

Field	Description
OSPF process	OSPF process identifier.
Destination	IP address of the destination port.
Metric	Number of tunnel metric.
Tunnel-ID	Tunnel address (destination port) for the session.
Tunnel-End-Point	Tunnel endpoint address of ospf.

show ip ospf interface

Use this command to display interface information for OSPF.

Command Syntax

```
show ip ospf interface (IFNAME|)
```

Parameters

IFNAME	Interface name.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip ospf interface
eth1 is up, line protocol is up
 Internet Address 10.100.10.72/24, Area 0.0.0.0, MTU 1500
    Router ID 100.100.100.72, Network Type BROADCAST, Cost: 10, TE Metric 0
    Transmit Delay is 1 sec, State DR, Priority 1
       Holddown timer : 50 seconds, Remaining time =  30seconds
    Designated Router (ID) 100.100.100.72, Interface Address 10.100.10.72
    Backup Designated Router (ID) 10.100.12.57, Interface Address 10.100.10.105
    Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
      Hello due in 00:00:05
    Neighbor Count is 1, Adjacent neighbor count is 1
    Crypt Sequence Number is 0
    Hello received 19 sent 106, DD received 4 sent 3
    LS-Req received 1 sent 1, LS-Upd received 3 sent 3
    LS-Ack received 2 sent 3, Discarded 0
```

OSPF Interface Fields

Table 1-25 explains the fields for each interface entry.

Table 1-25: OSPF interface output details

Field	Description
Internet address	IP address and subnet mask of the interface.
Area	OSPF area to which the interface belongs.
MTU	Maximum Transmission Unit (MTU) of the interface.
Transmit Delay	Transmit delay of the interface.

Table 1-25: OSPF interface output details

Field	Description
Priority	OSPF priority of the interface used for election of Designated Router (DR) and Backup Designated Router (BDR).
Hello	OSPF hello-interval.
Dead	OSPF dead-interval.
Wait	Hello wait-interval.
Retransmit	The period, in seconds, for which the router waits between retransmissions of OSPF packets that have not been acknowledged.
Hello due in	Time period for which router expects to receive hello packet.
Neighbor Count	OSPF neighbor count.
Adjacent neighbor	OSPF adjacent neighbor count.
Crypt Sequence Number	Used for authentication.
Hello received	Number of Hello packets and DD packets sent and received.
LS-Req	Number of LSA requests and LSA updates sent and received.
LS-Ack	Number of LSA acknowledgments sent and received number of LSA acknowledgment discards.

Example: DoNotAge

The following is sample output of this command when DoNotAge is enabled:

```
#show ip ospf interface eth1
eth1 is up, line protocol is up
 Internet Address 1.1.1.1/24, Area 0.0.0.0, MTU 1500
 Process ID 0, Router ID 33.33.33.33, Network Type BROADCAST, Cost: 10
 Transmit Delay is 1 sec, State Waiting, Priority 1, TE Metric 0
 No designated router on this network
 No backup designated router on this network
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02
 Neighbor Count is 0, Adjacent neighbor count is 0
 Crypt Sequence Number is 1106347721
 Hello received 0 sent 1, DD received 0 sent 0
 LS-Req received 0 sent 0, LS-Upd received 0 sent 0
 LS-Ack received 0 sent 0, Discarded 0
 Reduce LSA flooding
```

Example: Hello Suppression

The following is sample output of this command when Hello-Suppression is enabled:

```
#sh ip os interface
p7p1 is up, line protocol is up
  Internet Address 14.1.1.2/24, Area 0.0.0.0, MTU 1500
```

```
  Process ID 1, VRF (default), Router ID 2.2.2.2, Network Type POINTOMULTIPOINT, Cost:
1
  Reduce LSA flooding.
  Transmit Delay is 1 sec,  State Point-To-Point, TE Metric 1
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    Hello due in 00:00:03
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 1 neighbor(s)
  Hello received 5 sent 8, DD received 8 sent 6
  LS-Req received 2 sent 2, LS-Upd received 8 sent 9
  LS-Ack received 6 sent 6, Discarded 0
  No authentication
```

show ip ospf multi-area-adjacencies

Use this command to display multi-area adjacency information for OSPF.

Command Syntax

```
show ip ospf (<0-65535>|) multi-area-adjacencies
```

Parameters

<0-65535>	The ID of the router process for which information should be displayed.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of this command:

```
#show ip ospf 1 multi-area-adjacencies

Multi-area-adjacency on interface eth1 to neighbor 20.20.20.10
Internet Address 20.20.20.11/24, Area 0.0.0.1, MTU 1500
Process ID 1, Router ID 10.10.10.10, Network Type POINTOPOINT, Cost: 10
Transmit Delay is 1 sec, State Point-To-Point
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02
Neighbor Count is 0, Adjacent neighbor count is 0
Crypt Sequence Number is 1229928206
Hello received 0 sent 513, DD received 0 sent 0
LS-Req received 0 sent 0, LS-Upd received 0 sent 0
LS-Ack received 0 sent 0, Discarded 0
```

Table 1-26 explains the fields for each adjacency entry.

Table 1-26: show ip ospf multi-area-adjacencies output details

Field	Description
Multi-area-adjacency	Specifies the interface name and the router ID to which it is connected.
Internet Address	As Stated
Area	As Stated
MTU	Maximum Transmission Unit in bytes.
Process ID	The Process Identifier.
Router ID	As Stated
Network Type	In multi-area adjacencies, this is a point-to-point network with the neighbor.

Table 1-26: show ip ospf multi-area-adjacencies output details (Continued)

Field	Description
Cost	A reference bandwidth of 100 Mbps for cost calculation. The formula to calculate the cost is reference bandwidth divided by interface bandwidth.
Transmit Delay	A stated
State	As stated
Timer intervals configured	Hello timer = 10, Dead timer = 40, Wait timer = 40, Retransmit timer = 5
Hello due in	Countdown timer for a Hello message from the neighbor.
Neighbor Count	The number of neighbor.
Adjacent neighbor count	The number of neighbors participating in adjacencies.
Crypt Sequence Number	The 32-bit cryptographic sequence number appended on each OSPF protocol packet.
Hello received, sent	Hello packets sent and received.
DD received, sent	Database Description packets sent and received.
LS-Req received, sent	Link State Request packets sent and received.
LS-Upd received, sent	Link State Update packets sent and received.
LS-Ack received, sent, discarded	Link State Acknowledgement packets sent, received, or discarded.

show ip ospf neighbor

Use this command to display information about OSPF neighbors.

Command Syntax

```
show ip ospf (<0-65535>|) neighbor
show ip ospf (<0-65535>|) neighbor all
show ip ospf (<0-65535>|) neighbor interface A.B.C.D
show ip ospf (<0-65535>|) neighbor A.B.C.D
show ip ospf (<0-65535>|) neighbor A.B.C.D detail
show ip ospf (<0-65535>|) neighbor detail
show ip ospf (<0-65535>|) neighbor detail all
```

Parameters

`<0-65535>`	The ID of the router process
`all`	Include downstatus neighbor
`A.B.C.D`	IPv4 address
`detail`	Details of neighbors

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID     Pri   State       Dead Time   Address      Interface    Instance ID
1.1.1.1           1   Full/ -     inactive    14.1.1.1     p7p1                0
3.3.3.3           1   Full/ -     00:01:41    15.1.1.2     p8p1                0
3.3.3.3           1   Full/ -     inactive    15.1.1.2     VLINK0
```

OSPF Neighbor Fields

Table 1-27 explains the fields for each neighbor entry.

Table 1-27: OSPF neighbor output details

Field	Description
OSPF process	OSPF process identifier.
Neighbor ID	OSPF router identifier of the neighbor.

Table 1-27: OSPF neighbor output details

Field	Description
Pri	OSPF priority of the neighbor.
State	State of the OSPF neighbor: DependUpon: dummy state Down: no OSPF neighbors detected at this instant Attempt: in an NBMA environment, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval Init: hello packet received, but the receiving router's ID was not included in the hello packet 2-Way: bi-directional communication established between two routers ExStart: master and slave roles determined Exchange: database description packets (DBD) sent Loading: exchange of LSRs (link state request) and LSUs (link state update) packets Full: routers fully adjacent with each other.
Dead Time	If a new Hello is not received within this duration, the neighbor is declared dead.
Address	IP address of neighbor's interface attached to the network.
Interface	The interface attached to the network on which the neighbor is located.
Instance ID	Instance identifier for the session.

Example: Detail Parameter

This is sample output from the command when the `detail` parameter is specified:

```
#show ip ospf neighbor detail
 Neighbor 10.10.10.50, interface address 10.10.10.50
 In the area 0.0.0.0 via interface eth0
 Neighbor priority is 1, State is Full, 5 state changes
 DR is 10.10.10.50, BDR is 10.10.10.10
 Options is 0x42 (*|O|-|-|-|-|E|-)
 Dead timer due in 00:00:38
 Neighbor is up for 00:53:07
 Database Summary List 0
 Link State Request List 0
 Link State Retransmission List 0
 Crypt Sequence Number is 0
 Thread Inactivity Timer on
 Thread Database Description Retransmission off
```

OSPF Neighbor Detail Fields

Table 1-28 explains the fields for each neighbor detail entry.

Table 1-28: OSPF neighbor output detail

Field	Description
Neighbor	OSPF router identifier of the neighbor.
interface address	IP address of the neighbor interface.
In the area	Area and interface through which the OSPF neighbor is known.
Neighbor priority	OSPF priority of the neighbor.
State	OSPF state as explained in Table 1-27.
state changes	Number of state changes since the neighbor was created.
Hello is suppressed	Hello suppression is enabled.
Poll interval	Poll timer value.
DR is	Router ID of the designated router for the interface.
BDR is	Router ID of the backup designated router for the interface.
Options	LSA options as explained in Table 1-17.
LLS Options	LSDB Resynchronization (LR) Restart Signal (RS-bit) Whether link-local signalling (LLS) and out-of-band (OOB) link-state database resynchronization are performed for nonstop forwarding (NSF).
OOB-Resync in progress (receiver)/last OOB-Resync	Last successful OOB resynchronization with the NSF-capable neighbor. The router waits before taking a neighbor adjacency down if the OOB resynchronization has not taken place since the time a restart signal (Hello packet with RS-bit set) was received from the neighbor.
Dead timer due in	Expected time before declaring the neighbor dead.
Poll due in	Poll timer thread.
Neighbor is up for	Time since the neighbor went into the two-way state.
Database Summary List	Number of LSAs in the neighbor’s database.
Link State Request List	Number of LSAs that need to be received from this neighbor to synchronize the neighbors' topological databases.
Link State Retransmission List	Number of advertisements flooded out an adjacency. To ensure flooding is reliable, advertisements are retransmitted until they are acknowledged.
Crypt Sequence Number is	MD5 cryptographic sequence number.
Thread Inactivity Timer	Off if hello suppression is enabled, on otherwise.
Thread Database Description Retransmission	Off if hello suppression is enabled, on otherwise.
Thread Link State Request Retransmission	Off if hello suppression is enabled, on otherwise.

Table 1-28: OSPF neighbor output detail (Continued)

Field	Description
Thread Link State Update Retransmission	Off if hello suppression is enabled, on otherwise.
Thread Poll Timer	Whether the poll timer thread is on.
Bidirectional Forwarding Detection is enabled	Status of BFD, enabled or disabled.

Example: Hello-Suppression Option

This is sample output from the command when the `detail` parameter is specified and Hello-Suppression is enabled:

```
#sh ip os neighbor detail
 Neighbor 1.1.1.1, interface address 14.1.1.1
    In the area 0.0.0.0 via interface p7p1
    Neighbor priority is 1, State is Full, 5 state changes
    Hello is suppressed
    DR is 0.0.0.0, BDR is 0.0.0.0
    Options is 0x62 (-|O|DC|-|-|-|E|-)
    Dead timer due in inactive
    Neighbor is up for 00:05:03
    Database Summary List 0
    Link State Request List 0
    Link State Retransmission List 0
    Crypt Sequence Number is 0
    Thread Inactivity Timer off
    Thread Database Description Retransmission off
    Thread Link State Request Retransmission off
    Thread Link State Update Retransmission off
```

show ip ospf route

Use this command to display the OSPF routing table.

Command Syntax

```
show ip ospf (<0-65535>|) route ( A.B.C.D |A.B.C.D/M |summary |)
show ip ospf (<0-65535>|) route ( A.B.C.D |A.B.C.D/M |summary | fast-reroute |)
```

Parameters

`<0-65535>`	Router process identifier.
`A.B.C.D`	Single route.
`A.B.C.D/M`	Single exact match route.
`summary`	Route counts.
`fast-reroute`	Fast-reroute routes.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip ospf route
OSPF process 10:
Codes: C - connected, D - Discard, 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
C 50.0.0.0/24 [10] is directly connected, eth1, Area 0.0.0.10
C 60.0.0.0/24 [10] is directly connected, eth3, Area 0.0.0.10
OSPF process 15:
Codes: C - connected, D - Discard, 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
C 80.0.0.0/24 [1] is directly connected, eth4, Area 0.0.0.15
```

Header

Each entry in this table has a code preceding it indicating the source of the routing entry.

Table 1-29 explains the fields of route codes.

Table 1-29: route codes

Code	Meaning	Description
C	connected	Routes directly connected to the local device that were not distributed via IGP. The device inherently knows of these networks, so there is no need to learn about these from another device. Connected routes are preferred over routes for the same network learned from routing protocols.
O	OSPF	Modifiers: 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
D	discard	An ABR or ASBR performing summarization installs a discard route in its routing table for the summarized network range to prevent routing loops where portions of the summarized network range do not have a more specific route in the RIB. External and internal discard route entries are installed by default. During route summarization, routing loops can happen if data sent to a nonexisting network appears to be a part of the summary, and the router doing the summarization has a less specific route that points back to the sending router for the network.

Route Entry Fields

Table 1-30 shows the route entry fields.

Table 1-30: route entry output details

Field	Description
Codes	As explained in Table 1-29.
IP address	IP address of the remote network.
Metric	For OSPF the metric is cost, which indicates the best quality path to use to forward packets.
Next hop router IP address	This route is available through the next hop router located at this IP address. This identifies exactly where packets go when they match this route.
Outgoing interface name	Interface used to get to the next-hop address for this route.
Area	OSPF area identifier

Example: Process Identifier

The following is a sample output with the process identifier parameter.

```
#show ip ospf 10 route
OSPF process 10:
Codes: C - connected, D - Discard, 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
C 50.0.0.0/24 [10] is directly connected, eth1, Area 0.0.0.10
C 60.0.0.0/24 [10] is directly connected, eth3, Area 0.0.0.10
```

show ip ospf valid

Use this command to display information about opaque LSAs.

Command Syntax

```
show ip ospf (<0-65535>|) opaque-link valid
```

Parameters

`<0-65535>`	The ID of the router process for which information will be displayed.
`opaque-link`	Displays information about the opaque link-local LSAs.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip ospf 1 opaque-link valid
```

show ip ospf virtual-links

Use this command to display virtual link information.

Command Syntax

```
show ip ospf (<0-65535>|) virtual-links (brief|)
```

Parameters

`<0-65535>`	The ID of the router process for which information will be displayed.
`brief`	Display summary of OSPF virtual-links.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following is the display of the virtual link information for two routers, one with the virtual link up and one with virtual link down.

```
ospfd#show ip ospf virtual-links
Virtual Link VLINK0 to router 10.10.0.9 is up
 Transit area 0.0.0.1 via interface eth0
 Transmit Delay is 1 sec, State Point-To-Point,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02
 Adjacency state Full
Virtual Link VLINK1 to router 10.10.0.123 is down
 Transit area 0.0.0.1 via interface *
 Transmit Delay is 1 sec, State Down,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in inactive
 Adjacency state Down
```

The following is the display of the virtual link information for two routers, one with the virtual link up and one with virtual link down when flood reduction is enabled

```
ospfd#show ip ospf virtual-links
Virtual Link VLINK0 to router 10.10.0.9 is up
 Transit area 0.0.0.1 via interface eth0
 Transmit Delay is 1 sec, State Point-To-Point,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:02
 Adjacency state Full
Virtual Link VLINK1 to router 10.10.0.123 is down
 Transit area 0.0.0.1 via interface *
 Transmit Delay is 1 sec, State Down,
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in inactive
 Adjacency state Down
```

```
    DoNotAge LSA Allowed
```

If Hello-Suppression is enabled

```
    M1#sh ip os virtual-links
    Virtual Link VLINK0 to router 3.3.3.3 is up
      Transit area 0.0.0.1 via interface p8p1
      Hello suppression enabled
      DoNotAge LSA allowed
      Local address 15.1.1.1/32
      Remote address 15.1.1.2/32
      Transmit Delay is 1 sec, State Point-To-Point,
      Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
        Hello due in inactive
      No authentication
        Adjacency state Full
    M1#
```

Table 1-31 explains the fields for each virtual-links entry.

Table 1-31: show ip ospf virtual-links output details

Field	Description
Virtual Link	Virtual link name, the router ID to which it is connected, and the state of the link.
Transit area	Transit area ID, the interface it uses, and its instance ID – an Instance ID should default to 0. It is only necessary to assign a value other than 0 on those links that will contain multiple separate communities of OSPF routers.
Local address	The local IP address and subnet mask.
Remote address	The remote IP address and subnet mask.
Transmit Delay	The delay, in seconds, between link-state transmits. This value must be the same for all nodes on the network. The range is 0 to 65535. The default is 1. The state is point-to-point.
Timer intervals configured	The configured values in seconds of the following timers: Hello, Dead, Wait, Retransmit.
Hello due in	A countdown timer that indicates when the next Hello packet should arrive.
Adjacency State	Whether the adjacency state is either up or down.

show ip protocols

Use this command to display OSPF process parameters and statistics.

Command Syntax

```
show ip protocols
show ip protocols ospf
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This is an example of the output from the `show ip protocols` command:

```
#show ip protocols
Routing Protocol is "ospf 200"
  Invalid after 0 seconds, hold down 0, flushed after 0
  Outgoing update filter list for all interfaces is
 Redistributed kernel filtered by filter1
  Incoming update filter list for all interfaces is
  Redistributing: kernel
  Routing for Networks:
 192.30.30.0/24
 192.40.40.0/24
  Routing Information Sources:
 GatewayDistanceLast Update
  Distance: (default is 110)
 AddressMaskDistance List
```

Table 1-32 explains the fields for each ip protocol entry.

Table 1-32: show ip protocols output details

Field	Description
Routing Protocol is "ospf 200"	Specifies the routing protocol used.
Invalid after 0 seconds	Specifies the value of the invalid parameter.
Hold down 0	Specifies the current value of the hold-down parameter.
Flushed after 0	Specifies the time in seconds after which the individual routing information will be thrown (flushed) out.
Outgoing update	Specifies whether the outgoing filtering list has been set.

Table 1-32: show ip protocols output details

Field	Description
Incoming update	Specifies whether the incoming filtering list has been set.
Redistributing	Lists the protocol that is being redistributed.
Routing for Networks	Specifies the networks for which the routing process is currently injecting routes.
Routing Information Sources	Lists all the routing sources the IP Infusion software is using to build its routing table.

show ip route fast-reroute

Use this command to display routes with alternate next hops.

Command Syntax

```
show ip route fast-reroute
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip route fast-reroute
```

shutdown

Use the this command to temporarily shut down a protocol in the least disruptive manner and to notify its neighbors that it is going away.

Use the no parameter of this command.

Command Syntax

```
shutdown
no shutdown
```

Parameters

None

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#shutdown

#configure terminal
(config)#router ospf 100
(config-router)#no shutdown
```

snmp context-name

Use this command to configure the SNMP context name which will be mapped to each OSPFv2 instance.

When context name is configured OSPF OID will be registered with SNMP to that context.

Use `no` form of the command to un-map the SNMP context name from OSPFv2 instance and unregister OSPF OID from SNMP.

Note: snmp context-name will not be allowed in default ospf context.

Command Syntax

```
snmp context-name <WORD>
no snmp context-name <WORD>
```

Parameters

`WORD`	snmp context-name mapped to ospf instance. max length is 32

Command Mode

Router OSPF mode

Applicability

This command was introduced before OcNOS version 5.1.

Examples

```
(config)#router ospf 1
(config-router)#snmp context-name context1
```

snmp restart ospf

Use this command restart SNMP in OSPF

Command Syntax

```
snmp restart ospf
```

Parameter

None

Default

By default, SNMP resart is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#snmp restart ospf
```

summary-address

Use this command to summarize or suppress external routes with the specified address range.

Use the `no` option with this command to disable summary address.

An address range is a pairing of an address and a mask that is almost the same as IP network number. For example, if the specified address range is 192.168.0.0/255.255.240.0, it matches 192.168.1.0/24, 192.168.4.0/22, 192.168.8.128/25 and so on.

Redistributing routes from other protocols into OSPF requires the router to advertise each route individually in an external LSA. Use this command to advertise one summary route for all redistributed routes covered by a specified network address and mask. This minimizes the size of the OSPF link state database.

Command Syntax

```
summary-address (A.B.C.D/M | A.B.C.D A.B.C.D) (not-advertise|tag <0-4294967295>|)
no summary-address (A.B.C.D/M | A.B.C.D A.B.C.D) (not-advertise|tag)
```

Parameters

`A.B.C.D/M`	The range of addresses given as IPv4 starting address and a mask.
`A.B.C.D`	IP summary prefix e.g. i.i.i.i
`A.B.C.D`	IP summary prefix mask e.g. m.m.m.m
`not-advertise`	Suppress routes that match the range.
`tag`	Tag value to use as a “match” value for controlling redistribution via route maps.
`<0-4294967295>`	Set a tag value.

Default

By default, tag value is 0

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example uses the `summary-address` command to aggregate external LSAs that match the network 172.16.0.0/24 and assign a tag value of 3.

```
#configure terminal
(config)#router ospf 100
(config-router)#summary-address 172.16.0.0/16 tag 3
```

timers lsa arrival

This command sets the minimum interval to accept the same link-state advertisement (LSA) from OSPF neighbors.

Use the `no` form of this command to restore the default value.

Command Syntax

```
timers lsa arrival <0-600000>
no timers lsa arrival
```

Parameters

`<0-600000>`	The minimum delay in milliseconds between accepting the same LSA from neighbors.

Default

By default, Minimum LSA Arrival timer is 1 sec.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#timers lsa arrival 5000
```

timers spf exp

Use this command to set the Shortest-Path First (SPF) best-path schedule minimum and maximum delay between receiving a change to SPF calculation in milliseconds.

Use no parameter of this command to unset the SPF best-path schedule.

Command Syntax

```
timers spf exp <0-2147483647> <0-2147483647>
no timers spf exp
```

Parameters

<0-2147483647>	The minimum delay in milliseconds between receiving a change to SPF calculation.
<0-2147483647>	The maximum delay in milliseconds between receiving a change to SPF calculation.

Default

Default minimum delay: 500 milliseconds

Default maximum delay: 50000 milliseconds (50 seconds)

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#timers spf exp 300 300
```

timers throttle lsa

This command sets the rate-limiting intervals for OSPF link-state advertisement (LSA) generation.

Use the `no` form of this command to restore the default values.

Command Syntax

```
timers throttle lsa all <0-600000> <1-600000> <1-600000>
no timers throttle lsa all
```

Parameters

`<0-600000>`	Start interval: The minimum delay in milliseconds for the generation of LSAs. The first instance of LSA is always generated immediately upon a local OSPF topology change. The generation of the next LSA is not before the start interval.
`<0-600000>`	Hold interval: The hold time in milliseconds. This value is used to calculate the subsequent rate limiting times for LSA generation.
`<0-600000>`	Maximum interval: The maximum wait time in milliseconds between generation of the same LSA.

Defaults

Default start interval: 0 milliseconds

Default hold interval: 5000 milliseconds

Default maximum interval: 5000 milliseconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#timers throttle lsa all 200 10000 45000
```

CHAPTER 2 OSPFv2 Graceful Restart Commands

This chapter describes the OSPFv2 graceful restart commands.

- action interface
- capability restart
- event nbr
- debug ip ospf graceful-restart
- if-event-track ospf event
- ospf restart grace-period
- ospf restart helper
- restart ospf graceful

action interface

Configure action for an interface whose cost has to be updated on neighbor going down.

Use the no parameter with this command to remove the configs.

Command Syntax

```
action interface IFNAME cost <1-65535>
```

Parameters

`interface`	Apply action on interface
`IFNAME`	Interface's name
`cost`	Interface cost
`<1-65535>`	cost

Command Mode

ospf interface event mode (config-ospf-if-event)

Default

None

Applicability

This command was introduced before OcNOS version 6.3.0.

Example

```
(config-ospf-if-event)#action interface eth1 cost 100
(config-ospf-if-event)#commit
(config-ospf-if-event)#no action interface eth1 cost
(config-ospf-if-event)#commit
```

capability restart

Use this command to enable OSPF graceful restart or restart signaling. If a router is not restart-enabled, it cannot enter graceful restart mode and act as a helper.

Use the `no` parameter with this command to disable the features.

Command Syntax

```
capability restart graceful
```

```
no capability restart graceful
```

Parameters

None

Default

By default, OSPF graceful restart or restart signaling is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#capability restart graceful

(config)#router ospf 100
(config-router)#no capability restart graceful
```

event nbr

Configure events with OSPF neighbor id to monitor the state of OSPF neighbor in order to trigger events on any of the down states.

Use the `no` parameter with this command to remove the configuration.

Command Syntax

```
event nbr A.B.C.D down
```

Parameters

`nbr`	Configure neighbor sub-event
`A.B.C.D`	Destination IPv4 address
`down`	Event type down

Default

None

Command Mode

ospf interface event mode (config-ospf-if-event)

Applicability

This command was introduced before OcNOS version 6.3.0.

Example

```
(config-ospf-if-event)#event nbr 10.1.1.1 down
(config-ospf-if-event)#commit
(config-ospf-if-event)#no event nbr 10.1.1.1 down
(config-ospf-if-event)#commit
```

if-event-track ospf event

Use this command to configure an interface-track event with a valid event name with either matching all events or matching any of the events configured.

Use the no parameter with this command to remove the configs.

Command Syntax

if-event-track ospf event NAME match < all | any >

Parameters

`ospf`	OSPF specific commands
`event`	Configure the event
`NAME`	Event Name
`match`	Match criteria
`all`	Match criteria all
`any`	Match criteria any

Default

None

Command Mode

Config Mode

Applicability

This command was introduced before OcNOS version 6.3.0.

Example

```
(config)#if-event-track ospf event test match all
(config-ospf-if-event)#commit
(config-ospf-if-event)#exit
(config)#no if-event-track ospf event test
(config)#commit
```

debug ip ospf graceful-restart

Use this command to specify debugging option for OSPF graceful restart.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ip ospf graceful-restart (detail|terse|)
no debug ip ospf graceful-restart (detail|terse|)
```

Parameters

`detail`	Debug OSPF graceful restart detail information
`terse`	Debug OSPF graceful restart summary information

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf graceful-restart detail
```

ospf restart grace-period

Use this command to set the grace period for restarting the router.

If graceful restart is enabled, NSM is notified about the grace period. If the OSPF daemon unexpectedly shuts down, NSM sends this value to the OSPF daemon when it comes up again which uses this value to end the graceful state.

Use the `no` parameter with this command to revert to the default.

Command Syntax

```
ospf restart grace-period <2-1800>
no ospf restart grace-period
```

Parameters

`<2-1800>`	Grace period in seconds.

Default

The default grace period for restarting the OSPF router is 120 seconds.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ospf restart grace-period 250
```

ospf restart helper

Use this command to configure the helper behavior for graceful restart.

Use the `no` parameter with this command to revert to default.

Command Syntax

```
ospf restart helper max-grace-period <2-1800>
ospf restart helper never (router-id A.B.C.D|)
no ospf restart helper never
no ospf restart helper (never router-id (A.B.C.D | all) | max-grace-period|)
```

Parameters

`max-grace-period`	Help only if received grace-period is less than this value.
`<2-1800>`	Help only if received grace-period is less than this value.
`router-id`	Neighbor to never to act as helper.
`A.B.C.D`	Router ID of neighbor to never to act as helper.
`never`	Prevent the neighbor from entering helper mode.
`all`	All neighbors to never to act as helper.

Default

By default, router behave as helper. To disable it as helper, `ospf restart helper never` command should be configured. `ospf restart helper max-grace-period` – Max-grace-period to function as helper. If not configured, value will be the grace-period in restarting node.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ospf restart helper never router-id 1.1.1.1

#configure terminal
(config)#no ospf restart helper never router-id all
```

restart ospf graceful

Use this command to restart OSPF gracefully.

After this command is executed, the router immediately shuts down. NSM is notified that OSPF has shut down gracefully. NSM preserves routes installed by OSPF until the grace period expires.

Command Syntax

```
restart ospf graceful (grace-period <1-1800>|)
```

Parameters

`<1-1800>`	Grace period in seconds.

Default

By default, restart ospf graceful is disabled. Default value of `grace-period` is 120 seconds.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#restart ospf graceful grace-period 200

#restart ospf graceful
% Warning : OSPF process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
```

CHAPTER 3 OSPFv3 Commands

This chapter provides an alphabetized reference for each of the OSPFv3 commands. It includes the following commands:

- abr-type
- area default-cost
- area nssa
- area range
- area stub
- area virtual-link
- auto-cost reference bandwidth
- bfd all-interfaces
- capability restart
- clear ipv6 ospf process
- debug ipv6 ospf
- debug ipv6 ospf bfd
- debug ipv6 ospf events
- debug ipv6 ospf ifsm
- debug ipv6 ospf lfa
- debug ipv6 ospf lsa
- debug ipv6 ospf nfsm
- debug ipv6 ospf nsm
- debug ipv6 ospf packet
- debug ipv6 ospf retransmission
- debug ipv6 ospf rib
- debug ipv6 ospf route
- default-information originate
- default-metric
- distance
- distribute-list
- enable db-summary-opt
- fast-reroute keep-all-paths
- fast-reroute tie-break
- ipv6 ospf authentication
- ipv6 ospf bfd
- ipv6 ospf cost
- ipv6 ospf dead-interval
- ipv6 ospf display route single-line
- ipv6 ospf link-lsa-suppression

- ipv6 ospf mtu-ignore
- ipv6 ospf neighbor
- ipv6 ospf network
- ipv6 ospf priority
- ipv6 ospf restart grace-period
- ipv6 ospf restart helper
- ipv6 ospf restart planned-only
- ipv6 ospf retransmit-interval
- ipv6 ospf transmit-delay
- ipv6 router ospf
- ipv6 te-metric
- log-adjacency-changes
- max-concurrent-dd
- passive-interface
- redistribute
- restart ipv6 ospf graceful
- router-id
- router ipv6 ospf
- show debugging ipv6 ospf
- show ipv6 ospf
- show ipv6 ospf database
- show ipv6 ospf interface
- show ipv6 ospf neighbor
- show ipv6 ospf route
- show ipv6 route fast-reroute
- show ipv6 ospfv3 topology
- show ipv6 ospf virtual-links
- show ipv6 vrf
- shutdown
- snmp restart ospf6
- summary-address
- timers spf exp

© 2025 IP Infusion Inc. Proprietary

abr-type

Use this command to set an OSPFv3 Area Border Router (ABR) type.

Use the `no` parameter with this command to revert the ABR type to the default setting (`cisco`).

Specifying the ABR type allows better functioning in a multi-vendor environment. The ABR types are:

- Cisco (RFC 3509): A router is considered an ABR if it has more than one area actively attached and one of them is the backbone area.
- IBM (RFC 3509): A router is considered an ABR if it has more than one area actively attached and the backbone area is configured. In this case the configured backbone need not be actively connected.
- Standard (RFC 2328): A router is considered an ABR if it has more than one area actively attached to it.

Command Syntax

```
abr-type (cisco|ibm|standard)
no abr-type
```

Parameters

`cisco`	Specify an alternative ABR using Cisco implementation. This is the default ABR type.
`ibm`	Specify an alternative ABR using IBM implementation.
`standard`	Specify a standard ABR.

Default

By default, ABR type is Cisco

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#abr-type standard
```

area default-cost

Use this command to specify the cost for default summary route sent into a stub or a NSSA area. If an area is configured as a stub, the OSPFv3 router originates one type-3 inter-area-prefix-LSA into the stub area. This command changes the metric for this LSA.

Use the `no` parameter with this command to remove the assigned default cost.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) default-cost (<0-16777215>)
no area (A.B.C.D|<0-4294967295>) default-cost
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`<0-16777215>`	The advertised cost for the default summary route used for a stub or NSSA area.

Default

By default, advertised cost for the default summary route is 1.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#area 1 default-cost 10
```

area nssa

Use this command to set an area as a Not-So-Stubby-Area (NSSA). There are no external routes in an OSPF stub area, so you cannot redistribute from another protocol into a stub area. An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas. However, the external routes from other areas still do not enter the NSSA. You can configure an area to be a stub area or an NSSA, but not both.

This command simplifies administration when connecting a central site using OSPF to a remote site that is using a different routing protocol. You can extend OSPF to cover the remote connection by defining the area between the central router and the remote router as a NSSA.

Use the `no` form of this command to make an area a normal area.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) nssa
area (A.B.C.D|<0-4294967295>) nssa {translator-role (candidate|always)|stability-
  interval <0-2147483647>|no-redistribution|default-information-originate (metric
  <0-16777214>|metric-type <1-2>|)|no-summary}
no area (A.B.C.D|<0-4294967295>) nssa
no area (A.B.C.D|<0-4294967295>) nssa {translator-role|stability-interval|no-
  redistribution|default-information-originate|no-summary}
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`translator-role`	NSSA-ABR translator role:
`candidate`	Translate NSSA-LSA to Type-5 LSA if router is elected.
`always`	Always translate NSSA-LSA to Type-5 LSA.
`stability-interval`	Stability timer for a NSSA area. If an elected translator determines its services are no longer required, it continues to perform its duties for this time interval. This minimizes excess flushing of translated Type-7 LSAs and provides a more stable translator transition.
`<0-4294967295>`	Stability interval in seconds.
`no-redistribution`	Do not redistribute into the NSSA.
`default-information-originate`	Originate Type-7 default LSA into the NSSA.
`metric`	Specify metric for default routes.
`<0-16777214>`	Specify metric value.
`metric-type`	Specify metric type (see RFC 3101).

<1-2>	Specify metric type: 1: Type 1 external route 2: Type 2 external route
no-summary	Do not inject inter-area routes into the NSSA.
translate-candidate	Translate NSSA-LSA to Type-5 LSA if router is elected.
translate-always	Always translate NSSA-LSA to Type-5 LSA.

Default

By default, the nssa option value is candidate.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router ipv6 ospf
(config-router)#area 3 nssa translator-role candidate no-redistribution
default-information-originate metric 34 metric-type 2
```

area range

Use this command to configure the OSPF address range. This command summarizes intra-area routes for an area. The single summary route is then advertised to other areas by the Area Border Routers (ABRs). Routing information is condensed at area boundaries and outside the area. If the network numbers in an area are assigned in a way such that they are contiguous, the ABRs can be configured to advertise a summary route that covers all the individual networks within the area that fall into the specified range.

Use the `no` parameter with this command to remove the assigned area range.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) range X:X::X:X/M
area (A.B.C.D|<0-4294967295>) range X:X::X:X/M (not-advertise|)
no area (A.B.C.D|<0-4294967295>) range X:X::X:X/M (not-advertise|)
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`X:X::X:X/M`	The area IPv6 range prefix and length.
`not-advertise`	Does not advertise this range.

Default

By default, OSPF address range is advertised.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#area 1 range 2000::/3
```

area stub

Use this command to define an area as a stub area on all routers. There are two stub area router configuration commands: the `stub` and commands. In all routers attached to the stub area, configure the area by using the `stub` option of the area command. For an area border router (ABR) attached to the stub area, use the `area` command.

Use the `no` form of this command to make an area a normal area.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) stub
area (A.B.C.D|<0-4294967295>) stub no-summary
no area (A.B.C.D|<0-4294967295>) stub
no area (A.B.C.D|<0-4294967295>) stub no-summary
```

Parameters

`A.B.C.D`	OSPF Area ID in IPv4 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`no-summary`	Stops an ABR from sending summary link advertisements into the stub area.

Default

No stub area is defined.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#area 1 stub no-summary
```

area virtual-link

Use this command to configure a link between two backbone areas that are physically separated through other nonbackbone areas.

Use the `no` parameter with this command to remove the virtual link.

In OSPFv3, all non-backbone areas must be connected to a backbone area. If the connection to the backbone is lost, the virtual link repairs the connection.You can configure virtual links between any two backbone routers that have an interface to a common non-backbone area. The protocol treats these two routers joined by a virtual link as if they were connected by an unnumbered point-to-point network.

Configure the `hello-interval` to be the same for all routers attached to a common network. If the `hello-interval` is short, the router detects topological changes faster, but more routing traffic follows.

The `retransmit-interval` is the expected round-trip delay between any two routers in a network. Set the value to be greater than the expected round-trip delay to avoid needless retransmissions.

The `transmit-delay` is the time taken to transmit a link state update packet on the interface. Before transmission, the link state advertisements in the update packet are increased by this amount. Set the transmit-delay to be greater than zero. Also, take into account the transmission and propagation delays for the interface.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D fall-over bfd
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D (dead-interval <1-65535>|hello-
  interval <1-65535>|retransmit-interval <1-1800>|transmit-delay <1-1800>)
area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D instance-id (<0-31>|<64-95>)
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D fall-over bfd
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D (dead-interval|hello-
  interval|retransmit-interval|transmit-delay)
no area (A.B.C.D|<0-4294967295>) virtual-link A.B.C.D instance-id
```

Parameters

`A.B.C.D`	OSPF Area ID in IP64 address format.
`<0-4294967295>`	OSPF Area ID as a decimal value.
`A.B.C.D`	Specify router ID associated with a virtual link neighbor.
`dead-interval`	The interval in seconds during which no packets are received and after which the router acknowledges a neighboring router as off-line.
`<1-65535>`	The timer interval.
`hello-interval`	The interval in seconds the router waits before it sends a hello packet.
`<1-65535>`	The timer interval.
`retransmit-interval`	The interval in seconds the router waits before it retransmits a packet.
`<1-1800>`	The timer interval.
`transmit-delay`	The interval in seconds the router waits before it transmits a packet.

`<1-1800>`	The timer interval.
`instance-id`	The OSPFv3 instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.
`fall-over bfd`	Fall-over Bidirectional Forwarding Detection (BFD).

Default

Default hello interval:10 seconds.

Default dead interval:40 seconds.

Default retransmit interval: 5 seconds.

Default transmit delay: 1 second

Default instance-id: 0

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#area 1 virtual-link 10.10.11.50 hello 5 dead 10
(config-router)#area 1 virtual-link 10.10.11.50 instance-id 1
(config-router)#area 1 virtual-link 10.10.11.50 fall-over bfd
```

auto-cost reference bandwidth

Use this command to control how OSPFv3 calculates default metrics for the interface.

By default, OSPFv3 calculates the OSPFv3 metric for an interface by dividing the reference bandwidth by the interface bandwidth. The auto-cost command is used to differentiate high bandwidth links. For multiple links with high bandwidth, specify a larger reference bandwidth value to differentiate cost on those links.

Use the `no` form of this command to assign cost based only on the interface bandwidth.

Command Syntax

```
auto-cost reference-bandwidth <1-4294967>
no auto-cost reference-bandwidth
```

Parameters

`<1-4294967>`	The reference bandwidth in Mbps per second.

Default

By default, reference bandwidth is 100Mbps

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example changes the reference bandwidth to 1Gbps to change the Fast Ethernet interface cost from 1 to 10.

```
#configure terminal
(config)#router ipv6 ospf 1
(config-router)#auto-cost reference-bandwidth 1000

(config)#router ipv6 ospf 1
(config-router)#no auto-cost reference-bandwidth
```

bfd all-interfaces

Use this command to enable Bidirectional Forwarding Detection (BFD) on all interfaces.

Use the `no` form of this command to disable BFD.

Command Syntax

```
bfd all-interfaces
no bfd all-interfaces
```

Parameters

None

Default

By default, BFD is disabled on all OSPF enabled interfaces.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example changes the reference bandwidth to 1Gbps to change the Fast Ethernet interface cost from 1 to 10.

```
#configure terminal
(config)#router ipv6 ospf 1
(config-router)#bfd all-interfaces
```

capability restart

Use this command to enable OSPFv3 graceful restart capability. If a router is not restart-enabled, it cannot enter graceful restart mode and act as a helper.

Use the `no` parameter with this command to disable the feature.

Command Syntax

```
capability restart graceful
no capability restart
```

Parameter

None

Default

By default, capability restart graceful is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf 100
(config-router)#capability restart graceful

(config)#router ipv6 ospf 100
(config-router)#no capability restart
```

clear ipv6 ospf process

Use this command to clear and restart all OSPFv3 routing processes or a given OSPFv3 routing process.

Command Syntax

```
clear ipv6 ospf (WORD|) process
```

Parameters

WORD	OSPFv3 process tag.

Command Mode

Privileged Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ipv6 ospf Tag1 process
```

debug ipv6 ospf

Use this command to specify all debugging options for OSPFv3.

Use the `no` form of this command to disable the options.

Command Syntax

```
debug ipv6 ospf (all|bfd|events|ifsm|lsa|nfsm|nsm|packet|retransmission|rib|route|)
no debug ipv6 ospf
  (all|bfd|events|ifsm|lsa|nfsm|nsm|packet|retransmission|rib|route|)
no debug all ipv6 ospf
no debug all
```

Parameters

`all`	Enables all debugging information.
`bfd`	Debug OSPFv3 Bidirectional Forwarding Detection. (see debug ipv6 ospf bfd)
`events`	Debug OSPFv3 events (see debug ipv6 ospf events).
`ifsm`	Debug OSPFv3 Interface State Machines (see debug ipv6 ospf ifsm).
`lsa`	Debug OSPFv3 Link State Advertisements (see debug ipv6 ospf lsa).
`nfsm`	Debug OSPFv3 Neighbor State Machines (see debug ipv6 ospf nfsm).
`nsm`	Debug OSPFv3 NSM information (see debug ipv6 ospf nsm).
`packet`	Debug OSPFv3 packets (see debug ipv6 ospf packet).
`retransmission`	Debug OSPFv3 retransmission information. (see debug ipv6 ospf retransmission)
`rib`	Debug OSPFv3 Routing Information Base.(see debug ipv6 ospf rib)
`route`	Debug OSPFv3 route information (see debug ipv6 ospf route).

Command Mode

Privileged Exec and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf all
```

© 2025 IP Infusion Inc. Proprietary

debug ipv6 ospf bfd

Use this command to specify the debugging options for OSPFv3 Bidirectional Forwarding Detection

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf bfd
no debug ipv6 ospf bfd
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf bfd
```

debug ipv6 ospf events

Use this command to display debug information related to OSPF internal events. Use this command without parameters to turn on all the options.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf events {(abr|asbr|os|router|vlink|nssa)|}
no debug ipv6 ospf events {(abr|asbr|os|router|vlink|nssa)|}
```

Parameters

`abr`	Debug ABR events
`asbr`	Debug ASBR events
`os`	Debug OS interaction events
`router`	Debug other router events
`vlink`	Debug virtual link events
`nssa`	Debug NSSA events

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ipv6 ospf events abr
#debug ipv6 ospf events asbr
```

debug ipv6 ospf ifsm

Use this command to specify debugging options for OSPFv3 Interface Finite State Machine (IFSM) troubleshooting.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf ifsm ({events|status|timers}|)
no debug ipv6 ospf ifsm ({events|status|timers}|)
```

Parameters

`events`	Debug IFSM event information.
`status`	Debug IFSM status information.
`timers`	Debug IFSM timer information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf ifsm status
```

debug ipv6 ospf lfa

Use this command to specify the debugging options for OSPFv3 Loop-free Alternate path

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf lfa
no debug ipv6 ospf lfa
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf lfa
```

debug ipv6 ospf lsa

Use this command to specify the debugging options for OSPFv3 Link State Advertisements (LSAs).

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf lsa {(generate|flooding|install|maxage|refresh)|}
no debug ipv6 ospf lsa {(generate|flooding|install|maxage|refresh)|}
```

Parameters

`generate`	Debug LSA generation.
`flooding`	Debug LSA flooding.
`install`	Debug LSA installation.
`maxage`	Debug the maximum age processing.
`refresh`	Debug LSA refresh.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf lsa
```

debug ipv6 ospf nfsm

Use this command to specify debugging options for OSPFv3 Neighbor Finite State Machines (NFSMs).

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf nfsm {(events|status|timers)|}
no debug ipv6 ospf nfsm {(events|status|timers)|}
```

Parameters

`events`	Debug NFSM event information.
`status`	Debug NFSM status information.
`timers`	Debug NFSM timer information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf nfsm events
#no debug ipv6 ospf nfsm timers
```

debug ipv6 ospf nsm

Use this command to specify the debugging options for OSPFv3 NSM information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf nsm {(interface|redistribute)|}
no debug ipv6 ospf nsm {(interface|redistribute)|}
```

Parameters

`redistribute`	Debug redistribute.
`interface`	Debug the NSM interface.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf nsm interface
```

debug ipv6 ospf packet

Use this command to specify the packet debugging options for OSPFv3 information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf packet ({hello|dd|ls-request|ls-update|ls-ack|send|recv|detail}|)
no debug ipv6 ospf packet ({hello|dd|ls-request|ls-update|ls-ack|send|recv|
  detail}|)
```

Parameters

`hello`	Debug OSPFv3 hello.
`dd`	Debug OSPFv3 database description.
`ls-request`	Debug OSPFv3 link state request.
`ls-update`	Debug OSPFv3 link state update.
`ls-ack`	Debug OSPFv3 link state acknowledgment.
`send`	Debug packets sent
`recv`	Debug packets received.
`detail`	Debug detail information.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf packet ls-request
```

debug ipv6 ospf retransmission

Use this command to specify the debugging options for OSPFv3 retransmission information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf retransmission
no debug ipv6 ospf retransmission
```

Parameters

None

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf retransmission
```

debug ipv6 ospf rib

Use this command to specify the debugging options for OSPFv3 RIB information.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf rib {(interface|redistribute)|}
no debug ipv6 ospf rib {(interface|redistribute)|}
```

Parameters

`redistribute`	Debug redistribute.
`interface`	Debug the NSM interface.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf rib interface
```

debug ipv6 ospf route

Use this command to specify which route calculation to debug. Use this command without parameters to turn on all the options.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ipv6 ospf route {(ase|ia|install|spf)|}
no debug ipv6 ospf route {(ase|ia|install|spf)|}
```

Parameters

`ase`	Debug external route calculations.
`ia`	Debug inter-area route calculations.
`install`	Debug the route installation.
`spf`	Debug the SPF calculation.

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ipv6 ospf route
#debug ipv6 ospf route ia
```

default-information originate

Use this command to create a default external route into an OSPF routing domain.

The system acts like an Autonomous System Boundary Router (ASBR) when you use the `default-information originate` command to redistribute routes into an OSPF routing domain. An ASBR does not by default generate a default route into the OSPF routing domain.

When you give the `default-information originate` command, also specify a `route-map` to avoid a dependency on the default network in the routing table.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
default-information originate
default-information originate {metric <0-16777214>|metric-type (1|2)|route-map
  WORD|always}
no default-information originate
no default-information originate {metric|metric-type|route-map|always}
```

Parameters

`always`	Used to advertise the default route regardless of whether there is a default route.
`metric`	Sets the OSPF metric used in creating the default route.
`<0-16777214>`	Sets the OSPF metric used in creating the default route. The value used is specific to the protocol.
`metric-type`	The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101).
`1`	Sets OSPF External Type 1 metric.
`2`	Sets OSPF External Type 2 metric (default).
`route-map`	Route map.
`WORD`	Specify the name of route map.

Default

Sets the OSPF metric used in creating the default route. The default metric value is 20. The value used is specific to the protocol. metric-type The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101).

By default, 2 sets OSPF External Type 2 metric

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#default-information originate always metric 23 metric-type 2
route-map myinfo

(config)#router ipv6 ospf
(config-router)#no default-information originate metric metric-type route-map
```

default-metric

Use this command to set a default metric for OSPF.

A default metric facilitates redistributing routes with incompatible metrics. If the metrics do not convert, the default metric provides an alternative. Use this command to use the same metric value for all redistributed routes. Use this command in conjunction with the redistribute command.

Use the `no` parameter with this command to return to the default state.

Command Syntax

```
default-metric <1-16777214>
no default-metric
```

Parameter

`<1-16777214>`	Default metric value.

Default

Default metric for redistributed routes will be 20.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#default-metric 100
```

distance

Use this command to define OSPFv3 route administrative distances based on route type. This command sets the distance for an entire group of routes rather than a specific route that passes an access list.

The administrative distance rates the trustworthiness of a routing information source. A higher distance value means a lower trust rating. For example, an administrative distance of 254 means that the routing information source cannot be trusted and should be ignored.

Use the `no` form of this command to restore the default value.

Command Syntax

```
distance <1-254>
distance ospfv3 {intra-area <1-254>|inter-area <1-254>|external <1-254>}
no distance (<1-254>|)
no distance ospfv3
```

Parameters

`<1-254>`	Used alone, this parameter specifies a default administrative distance used when no other specification exists for a routing information source.
`intra-area`	Routes within an area.
`<1-254>`	Distance for all routes within an area
`inter-area`	Routes from one area to another area.
`<1-254>`	Distance for all routes from one area to another area.
`external`	Routes from other routing domains learned by redistribution.
`<1-254>`	Distance for routes from other routing domains learned by redistribution.

Default

By default, distance value for each type of route is 110

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 ospf 100
(config-router)#distance ospfv3 inter-area 20 intra-area 10 external 40
```

distribute-list

Use this command to filter networks in routing updates. This command redistributes other routing protocols into the OSPF routing table.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
distribute-list WORD out ((kernel|connected|static|rip|bgp|isis|ospf (WORD|<1-
  65535>|)))
distribute-list WORD in
no distribute-list WORD out ((kernel|connected|static|rip|bgp|isis|ospf (WORD|<1-
  65535>|)))
no distribute-list WORD in
```

Parameters

`WORD`	Specify the name of the access list.
`in`	Filter incoming routing updates.
`out`	Filter outgoing routing updates.
`kernel`	Specify kernel routes.
`connected`	Specify connected routes.
`static`	Specify static routes.
`rip`	Specify RIP routes.
`bgp`	Specify BGP routes.
`isis`	Specify IS-IS routes.
`ospf`	Specify OSPF routes.
`WORD`	Specify the OSPF process tag. If not specified, redistribute OSPF process with tag "null".
`<1-65535>`	Specify OSPF process ID <1-65535>. If not specified, redistribute OSPF instance with process ID 0.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the distribution of BGP routing updates based on the access list `list1` (network 172.10.0.0).

```
#configure terminal
(config)#access-list list1 permit 172.10.0.0/16
```

```
(config)#router ipv6 ospf 100
(config-router)#distribute-list list1 out bgp
(config-router)#redistribute bgp
```

enable db-summary-opt

Use this command to enable the database summary list optimization for OSPFv3.

When this feature is enabled, the database exchange process is optimized by removing the LSA from the database summary list for the neighbor if the LSA instance in the summary list is the same as or less recent than the LSA in the database description packet received from the neighbor.

Use the `no` form of the command to disable database summary list optimization.

Command Syntax

```
enable db-summary-opt
```

```
no enable db-summary-opt
```

Parameters

None

Default

By default, db summary opt is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#enable db-summary-opt
(config-router)#no enable db-summary-opt
```

fast-reroute keep-all-paths

Use this command to enable fast rerouting on all OSPFv3 interfaces.

Use the `no` parameter with this command to disable fast rerouting.

Command Syntax

```
fast-reroute keep-all-paths
no fast-reroute keep-all-paths
```

Parameters

None

Defaults

By default, fast rerouting is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf 200
(config-router)#fast-reroute keep-all-paths
```

fast-reroute tie-break

Use this command to set the tie-breaking policy for selecting a fast reroute repair path. You assign a priority to each type of repair path.

Use the `no` parameter with this command to set the tie-breaking policy for a specific type of repair path to its default priority. To set all types of repair paths to their default priorities, do not specify a repair path with the `no` form of this command.

Command Syntax

```
fast-reroute tie-break (primary-path|interface-disjoint|node-protecting|broadcast-
  interface-disjoint) index <1-4>
no fast-reroute tie-break (primary-path|interface-disjoint|node-
  protecting|broadcast-interface-disjoint) index <1-4>
no fast-reroute tie-break
```

Parameters

`primary-path`	When there are multiple loop-free alternate paths. The primary path is selected for an frr if it is node-protecting as well as link-protecting
`interface-disjoint`	Do not select point-to-point interfaces that have no alternate next hop for rerouting if the primary gateway fails, thus protecting the interface.
`node-protecting`	Bypass the `primary-path` gateway router which might not protect the router that is the next hop in the primary path.
`broadcast-interface-disjoint`	Do not use the interface if connected to a broadcast network. Repair paths protect links when a repair path and a protected primary path use *different* next-hop interfaces. However, on broadcast interfaces, if the repair path is computed via the same interface as the primary path, but their next-hop gateways are different, the router is protected but the link might not be.
`index`	Tie break priority. A lower value has higher preference.
`<1-4>`	Range of priority values.

Defaults

The default priority scheme is:

1. primary-path
2. interface-disjoint
3. node-protecting
4. broadcast-interface-disjoint

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf 200
(config-router)#fast-reroute tie-break interface-disjoint index 1
```

ipv6 ospf authentication

Use this command to enable the authentication and/or confidentiality for OSPFv3 sessions on this interface using crypto map configuration.

Use no form of this command to disable the authentication and/or confidentiality for OSPFv3 sessions.

Command Syntax

```
ipv6 ospf authentication cryptomap WORD
no ipv6 ospf authentication cryptomap WORD
```

Parameters

`authentication`	IPsec Authentication
`cryptomap`	Map used to setup IPsec SA
`WORD`	Name of the Crypto-map

Default

No default value.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.0.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ipv6 ospf authentication cryptomap map1
```

ipv6 ospf bfd

Use this command to enable Bidirectional Forwarding Detection (BFD).

Use this command with either the no or disable parameter to disable BFD.

Command Syntax

```
ipv6 ospf bfd (disable|)
ipv6 ospf bfd (disable|)instance-id (<0-31>|<64-95>)
no ipv6 ospf bfd (disable|)
no ipv6 ospf bfd (disable|)instance-id (<0-31>|<64-95>)
```

Parameters

`disable`	Specify to disable BFD.
`instance-id`	Specify the instance. Default value for instance-id is 0.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast

Default

By default, IPv6 OSPF BFD is disabled. Default value for instance-id is 0.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ipv6 ospf bfd instance-id 2
```

Note: Interface level CLI: "ipv6 ospf bfd" will be displayed as "ipv6 ospf bfd instance-id 0" from 5.0 release.

ipv6 ospf cost

Use this command to specify the link-cost described in LSAs.

The cost (or metric) of an interface in OSPF indicates the overhead required to send packets across a certain interface. The value is taken to describe Link State information, and used for route calculation.

Use the `no` parameter with this command to reset the cost to default.

Command Syntax

```
ipv6 ospf cost <1-65535>
ipv6 ospf cost <1-65535> instance-id (<0-31>|<64-95>)
no ipv6 ospf cost
no ipv6 ospf cost instance-id (<0-31>|<64-95>)
```

Parameters

`cost`	Specify the link-state metric.
`<1-65535>`	Specify the link-state metric.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, ipv6 cost value is 10.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf cost 20 instance-id 1
```

ipv6 ospf dead-interval

Use this command to set the amount of time that the router waits to receive an OSPF hello packet from the neighbor before declaring the neighbor down.

The dead interval is advertised in hello packets. OSPF compares the dead interval in a received packet to the dead interval configured for the receiving interface. If the intervals do not match, the hello packet is discarded.

Use the `no` parameter with this command to reset the interval to default.

Command Syntax

```
ipv6 ospf dead-interval <1-65535>
ipv6 ospf dead-interval <1-65535> instance-id (<0-31>|<64-95>)
no ipv6 ospf dead-interval
no ipv6 ospf dead-interval instance-id (<0-31>|<64-95>)
```

Parameters

`dead-interval`	Specify the interval.
`<1-65535>`	Specify the interval in seconds.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, dead interval is 40 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf dead-interval 20
```

ipv6 ospf display route single-line

Use this command to display the output of the show ipv6 ospf route command with each route entry in a single-line.

Use the `no` parameter with this command to display the output with each route entry in a multiple lines.

Command Syntax

```
ipv6 ospf display route single-line
no ipv6 ospf display route single-line
```

Parameters

None

Default

By default, show ipv6 ospf route displays routes in multiple lines

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 ospf display route single-line
```

ipv6 ospf hello-interval

Use this command to specify the interval between `hello` packets.

The hello interval is advertised in the hello packets. An OSPF router compares the hello interval in a received packet to the interval configured for the receiving interface. If this interval does not match, the hello packet is discarded. A shorter hello interval ensures faster detection of topological changes, but results in more routing traffic.

Use the `no` parameter with this command to reset the interval to default.

Command Syntax

```
ipv6 ospf hello-interval <1-65535>
ipv6 ospf hello-interval <1-65535> instance-id (<0-31>|<64-95>)
no ipv6 ospf hello-interval
no ipv6 ospf hello-interval instance-id (<0-31>|<64-95>)
```

Parameters

`hello-interval`	Specify the interval.
`<1-65535>`	Specify the interval in seconds.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, hello interval is 10 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf hello-interval 5 instance-id 1
```

ipv6 ospf link-lsa-suppression

Use this command to enable or disable link LSA (type 8) suppression. A type 8 LSA gives information about link-local addresses and a list of IPv6 addresses on the link.

If enabled and the interface type is *not* broadcast or NBMA, the router does not send type 8 link LSAs. This implies that other routers on the link determine the router's next-hop address using a mechanism other than the type 8 link LSA. This feature is implicitly disabled if the interface type is broadcast or NBMA.

Command Syntax

```
ipv6 ospf link-lsa-suppression (enable|disable)
ipv6 ospf link-lsa-suppression (enable|disable) instance-id (<0-31>|<64-95>)
```

Parameters

`enable`	Enable type 8 link LSA suppression
`disable`	Disable type 8 link LSA suppression (default).
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, type 8 link LSA suppression is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf link-lsa-suppression enable
```

ipv6 ospf mtu-ignore

Use this command to configure OSPF so that it does not check the MTU size during DD (Database Description) exchange.

Use the `no` form of this command to make OSPF check the MTU size during DD exchange.

Command syntax

```
ipv6 ospf mtu-ignore
ipv6 ospf mtu-ignore instance-id (<0-31>|<64-95>)
no ipv6 ospf mtu-ignore
no ipv6 ospf mtu-ignore instance-id (<0-31>|<64-95>)
```

Parameters

`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, during the DD exchange process, OSPF checks the MTU size described in DD packets received from its neighbor. If the MTU size does not match the interface MTU, the neighbor adjacency is not established. Using this command makes OSPF ignore this check and allows establishing of adjacency regardless of MTU size in the DD packet.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#interface eth1
(config-if)#ipv6 ospf mtu-ignore
```

ipv6 ospf neighbor

Use this command to connect OSPFv3 routers to non-broadcast multi-access (NBMA) networks.

One neighbor entry must be included for each known NBMA neighbor. The neighbor address must be a link-local address.

Note: For point-to-multipoint interfaces, the `cost` parameter is the only applicable option.

Use the `no` parameter with this command to remove a configuration.

Command Syntax

```
ipv6 ospf neighbor X:X::X:X (instance-id (<0-31>|<64-95>)|)
ipv6 ospf neighbor X:X::X:X {cost <1-65535>} (instance-id (<0-31>|<64-95>)|)
ipv6 ospf neighbor X:X::X:X {poll-interval <0-4294967295>|priority <0-255>}
  (instance-id (<0-31>|<64-95>)|)
no ipv6 ospf neighbor X:X::X:X ({cost|poll-interval|priority}|) (instance-id (<0-
  31>|<64-95>)|)
```

Parameters

`X:X::X:X`	Specify a neighbor IP address.
`instance-id`	Specify the instance.
`<0-255>`	Specify the instance ID.
`cost`	Cost of the interface. This parameter does not apply to NBMA networks.
`<1-65535>`	Cost of the interface.
`poll-interval`	Dead neighbor polling interval.
`<0-4294967295>`	Dead neighbor polling interval in seconds. It is recommended to set this value much higher than the hello interval.
`priority`	Specify a priority. This parameter does not apply to point-to-multipoint interfaces.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

Default cost is 10.

Default poll interval is 120 seconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
```

```
(config)#interface eth0
(config-if)#ipv6 ospf neighbor 2000:500::1 cost 2 instance-id 3
```

ipv6 ospf network

Use this command to set an OSPFv3 network type.

Use the `no` option with this command to return to the default value.

Command Syntax

```
ipv6 ospf network (broadcast|non-broadcast|point-to-multipoint (non-
  broadcast|)|point-to-point) (instance-id (<0-31>|<64-95>)|)
no ipv6 ospf network (instance-id (<0-31>|<64-95>)|)
```

Parameters

`broadcast`	Sets the network type to broadcast.
`non-broadcast`	Sets the network type to NBMA.
`point-to-multipoint`	Sets the network type to point-to-multipoint.
`non-broadcast`	Sets the network type to NBMA.
`point-to-point`	Sets the network type to point-to-point.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, ipv6 ospf network is broadcast type

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows how to set the network to point-to-point type on the eth0 interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf network point-to-point
```

ipv6 ospf priority

Use this command to set the router priority for determining the designated router (DR) for the network.

A router with the higher router priority becomes the DR. If the priority is the same for two routers, the router with the higher router ID takes precedence.

Only routers with a nonzero priority value are eligible to become the designated or backup designated router. Configure router priority for broadcast or NBMA networks only and not for point-to-point networks.

Use the `no` parameter with this command to reset the value to default.

Command Syntax

```
ipv6 ospf priority <0-255>
ipv6 ospf priority <0-255> instance-id (<0-31>|<64-95>)
no ipv6 ospf priority instance-id (<0-31>|<64-95>)
```

Parameters

`priority`	Specify the router priority of the interface.
`<0-255>`	Specify the router priority of the interface. The default is 1.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, priority is 1

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf priority 127
```

ipv6 ospf restart grace-period

Use this command to enable the graceful restart feature and set the grace period for restarting the router.

If graceful restart is enabled, NSM is notified about the grace period. If the OSPF daemon unexpectedly shuts down, NSM sends this value to the OSPF daemon when it comes up again which uses this value to end the graceful state.

Use the `no` parameter with this command to revert to the default grace period.

Command Syntax

```
ipv6 ospf restart grace-period <2-1800>
no ipv6 ospf restart grace-period
```

Parameters

`grace-period`	Specify the grace period.
`<2-1800>`	Specify the grace period in seconds.

Default

By default, grace period is 120 seconds

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 ospf restart grace-period 250
```

ipv6 ospf restart helper

Use this command to configure the helper behavior for graceful restart.

Use the `no` parameter with this command to revert to the default.

Command Syntax

```
ipv6 ospf restart helper {max-grace-period <2-1800>}
ipv6 ospf restart helper never (router-id A.B.C.D|)
no ipv6 ospf restart helper
no ipv6 ospf restart helper never
no ipv6 ospf restart helper {max-grace-period|never router-id (A.B.C.D|all)}
```

Parameters

`<2-1800>`	Help only if received grace-period is less than this value.
`A.B.C.D`	Router ID of neighbor to never to act as helper.
`never`	Prevent the neighbor from entering helper mode.
`max-grace-period`	Help only if received grace-period is less than this value.
`router-id`	Router of neighbor to never to act as helper.

Default

By default, router behave as helper. To disable it as helper, `ospf restart helper never` command should be configured. `ospf restart helper max-grace-period` – Max-grace-period to function as helper. If not configured, value will be the grace-period in restarting node.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 ospf restart helper never router-id 1.1.1.1

#configure terminal
(config)#no ipv6 ospf restart helper never
```

ipv6 ospf restart planned-only

Use this command to configure the OSPFv3 GR support for only planned restart.

Use the `no` parameter with this command to revert to the default.

Note: OSPFv3 supports only planned restart. Hence this command configuration and un-configuration does not have any functionality impact

Command Syntax

```
ipv6 ospf restart planned-only
```

```
no ipv6 ospf restart planned-only
```

Parameters

None

Default

By default, OSPFv3 supports planned restart only.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#ipv6 ospf restart planned-only

#configure terminal
(config)#no ipv6 ospf restart planned-only
```

ipv6 ospf retransmit-interval

Use this command to set the interval between retransmission of Link State Update packets. This interval is also used to retransmit DD packets and Link State Request packets.

After sending an LSA to a neighbor, the router keeps the LSA on the LS-retransmission list until it receives an acknowledgement. If the router does not receive an acknowledgment from the neighbor during the retransmit interval, it sends the LSA to the neighbor again.

Use the `no` parameter with this command to reset the interval to the default value.

Command Syntax

```
ipv6 ospf retransmit-interval <1-1800>
ipv6 ospf retransmit-interval <1-1800> instance-id (<0-31>|<64-95>)
no ipv6 ospf retransmit-interval
no ipv6 ospf retransmit-interval instance-id (<0-31>|<64-95>)
```

Parameters

`retransmit-interval`	
	Specify the interval.
`<1-1800>`	Specify the interval in seconds.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, ipv6 ospf retransmit interval is 5 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf retransmit-interval 3
```

ipv6 ospf transmit-delay

Use this command to set the estimated time it takes to transmit a Link State Update packet over the interface. The transmit-delay value is added to the LS age of LSAs and is advertised through this interface whenever the LSAs are transmitted.

Use the `no` parameter with this command to reset the delay to the default value.

Command Syntax

```
ipv6 ospf transmit-delay <1-1800>
ipv6 ospf transmit-delay <1-1800> instance-id (<0-31>|<64-95>)
no ipv6 ospf transmit-delay
no ipv6 ospf transmit-delay instance-id (<0-31>|<64-95>)
```

Parameters

`transmit-delay`	Specify the time to transmit a link-state update.
`<1-1800>`	Specify the time in seconds to transmit a link-state update.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, transmit delay is 1 second

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf transmit-delay 2
```

ipv6 router ospf

Use this command to enable OSPFv3 routing on an interface.

Specify the process ID to configure multiple instances of OSPFv3. When running a single instance of OSPFv3, you do not need to specify a instance ID.

When OSPFv3 receives a packet, it checks if the instance ID in the packet matches the instance ID of the receiving interface.

Use the `no` parameter with this command to disable OSPFv3 routing on an interface.

Command Syntax

```
ipv6 router ospf area (A.B.C.D|<0-4294967295>)
ipv6 router ospf area (A.B.C.D|<0-4294967295>) instance-id (<0-31>|<64-95>)
ipv6 router ospf area (A.B.C.D|<0-4294967295>) tag WORD
ipv6 router ospf area (A.B.C.D|<0-4294967295>) tag WORD instance-id (<0-31>|<64-95>)
ipv6 router ospf tag WORD area (A.B.C.D|<0-4294967295>)
ipv6 router ospf tag WORD area (A.B.C.D|<0-4294967295>) instance-id (<0-31>|<64-95>)
no ipv6 router ospf area (A.B.C.D|<0-4294967295>)
no ipv6 router ospf area (A.B.C.D|<0-4294967295>) instance-id (<0-31>|<64-95>)
no ipv6 router ospf area (A.B.C.D|<0-4294967295>) tag WORD
no ipv6 router ospf area (A.B.C.D|<0-4294967295>) tag WORD instance-id (<0-31>|<64-95>)
no ipv6 router ospf tag WORD area (A.B.C.D|<0-4294967295>)
no ipv6 router ospf tag WORD area (A.B.C.D|<0-4294967295>) instance-id (<0-31>|<64-95>)
```

Parameters

`area`	OSPF Area ID in IPv4 address format.
`A.B.C.D`	OSPF area ID in IP address format.
`<0-4294967295>`	OSPF area ID as a decimal value.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.
`tag`	Tag value to use as a “match” value for controlling redistribution via route maps.
`WORD`	Set the tag value.

Default

By default, ipv6 router ospf is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 router ospf area 1 tag Tag1 instance-id 1
```

ipv6 te-metric

Use this command to set the traffic engineering metric for an interface.

The traffic engineering metric is used in OSPF-TE Link State Advertisements. If the traffic engineering metric is not set, the ipv6 ospf cost value for an interface is used in OSPF-TE Link State Advertisements.

Use the `no` parameter with this command to unset the traffic engineering metric for this interface.

Command Syntax

```
ipv6 te-metric <1-65535>
ipv6 te-metric <1-65535> instance-id (<0-31>|<64-95>)
no ipv6 te-metric
no ipv6 te-metric instance-id (<0-31>|<64-95>)
```

Parameters

`te-metric`	Specify the TE metric.
`<1-65535>`	Specify the TE metric value.
`instance-id`	Specify the instance.
`<0-31>`	Interface instance ID for IPv6 unicast
`<64-95>`	Interface instance ID for IPv4 unicast.

Default

By default, traffic engineering metric value is 0

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 te-metric 6
```

log-adjacency-changes

Use this command for the router to send a SYSLOG message when an OSPF neighbor goes up or down.

Use `no` parameter of this command to stop sending SYSLOG message.

Command Syntax

```
log-adjacency-changes (brief|detail|)
no log-adjacency-changes
```

Parameters

`detail`	Sends a SYSLOG message for each state change, not just when a neighbor goes up or down.

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#log-adjacency-changes brief
(config-router)#log-adjacency-changes detail

#configure terminal
(config)#router ipv6 ospf
(config-router)#no log-adjacency-changes
```

max-concurrent-dd

Use this command to limit the number of Database Descriptors (DD) that can be processed concurrently.

This command is useful when a router's performance is affected from simultaneously bringing up several OSPFv3 adjacencies. This command limits the maximum number of DD exchanges that can occur concurrently per OSPFv3 instance, thus allowing for all of the adjacencies to come up.

Use the `no` option with this command to remove the limit.

Command Syntax

```
max-concurrent-dd <1-65535>
no max-concurrent-dd
```

Parameters

`<1-65535>`	Specify the number of DD processes.

Default

By default, number of maximum concurrent DD processes is 5

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example set the `max-concurrent-dd` value to 4.

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#max-concurrent-dd 4
```

passive-interface

Use this command to suppress sending Hello packets on all interfaces, or on a specified interface.

This command configures OSPFv3 on simplex Ethernet interfaces. Since a simplex interface represents only one network segment between two devices, configure the transmitting interface as a passive interface. This ensures that OSPFv3 does not send hello packets for the transmitting interface. Both the devices can see each other via the hello packet generated for the receiving interface.

Use the `no` form with this command to resume sending `hello` packets on all interfaces, or on a specified interface.

Command Syntax

```
passive-interface
passive-interface IFNAME (disable|enable)
no passive-interface
no passive-interface IFNAME
```

Parameters

`IFNAME`	Specify an interface name

Default

By default, passive interface is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#passive-interface
(config-router)#passive-interface eth0 disable
```

redistribute

Use this command to import routes from other routing protocols, or from another OSPF instance, into OSPFv3 AS-external-LSAs.

OSPFv3 advertises routes learned from other routing protocols or from other OSPF instances, including static or connected routes. Each injected prefix is put into the AS-external-LSA with a specified metric and metric-type.

Use the `no` parameter with this command to stop redistribution.

Command Syntax

```
redistribute (connected|static|rip|bgp|isis|ospf (WORD|<1-65535>|)) {metric <0-16777214>|metric-type (1|2)|route-map WORD|tag <0-4294967295>}
no redistribute (connected|static|rip|bgp|isis|ospf (WORD|)) (metric|metric-type|route-map|tag|)
```

Parameters

`connected`	Specify connected routes.
`static`	Specify static routes.
`rip`	Specify RIP routes.
`bgp`	Specify BGP routes.
`isis`	Specify IS-IS routes.
`ospf`	Specify OSPF routes.
`WORD`	Specify an OSPFv3 Process Tag. If not specified, redistribute OSPF process with tag "null".
`<1-65535>`	Specify an OSPF process identifier. If not specified, redistribute OSPF instance with process ID 0.
`metric`	Specify the external metric.
`<0-16777214>`	Specify the external metric.
`metric-type`	Specify the external metric-type (see RFC 3101):
`1`	Set OSPF External Type 1 metric.
`2`	Set OSPF External Type 2 metric.
`route-map`	Specify a route map reference.
`WORD`	Specify name of the route-map.
`tag`	Tag value to use as a “match” value for controlling redistribution via route maps
`<0-4294967295>`	Specify the route tag.

Default

By default, redistribute is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows redistribution of BGP routes into the OSPFv3 routing table, with the metric as 10.

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#redistribute bgp metric 10 metric-type 1
```

The following example shows redistribution of static IPv4 routes into the OSPFv3 routing table.

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#redistribute static
```

restart ipv6 ospf graceful

Use this command to restart OSPFv3 gracefully.

After this command is executed, the router immediately shuts down. NSM is notified that OSPF has shut down gracefully. NSM preserves routes installed by OSPF until the grace period expires.

Command Syntax

```
restart ipv6 ospf graceful (grace-period <1-1800>|)
```

Parameters

grace-period	Specify a grace period.
<1-1800>	Specify a grace period in seconds.

Default

By default, restart ipv6 ospf graceful is disabled. Default value of grace-period is 120 seconds.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#restart ipv6 ospf graceful grace-period 200

#restart ipv6 ospf graceful
% Warning : ipv6 OSPF process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
```

router-id

Use this command to specify a router ID for the OSPFv3 process.

Configure each router with a unique router-id. In an OSPFv3 router process that has active neighbors, a new router-id is used at the next reload or when you start the OSPFv3 manually.

Use the `no` form of this command to force OSPFv3 to stop the routing functionality.

Command Syntax

```
router-id A.B.C.D
```

```
no router-id
```

Parameters

`A.B.C.D`	Specify the router ID in IPv4 address format.

Default

By default, router id is loop-back IP address of IP address with highest IP

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows a fixed router ID `43.3.3.3`

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#router-id 43.3.3.3
```

router ipv6 ospf

Use this command to initiate OSPFv3 routing process and enter Router mode to configure OSPFv3 routing process. For making the OSPFv3 routing process functional, you must specify OSPFv3 process tag in router mode and enable OSPFv3 on at least one interface. OSPFv3 is only enabled on interfaces where OSPFv3 process tag matches the tag specified using ipv6 router ospf area command in Interface mode.

Use the `no` parameter with this command to remove OSPFv3 process.

Command Syntax

```
router ipv6 ospf
router ipv6 ospf WORD
router ipv6 vrf ospf WORD
no router ipv6 ospf
no router ipv6 ospf WORD
no router ipv6 vrf ospf WORD
```

Parameters

`WORD`	Tag value to use as a “match” value for controlling redistribution via route maps.
`vrf`	Enable an IPv6 VRF routing process

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf Tag1
(config-router)#
```

show debugging ipv6 ospf

Use this command to display the OSPFv3 debugging options.

Command Syntax

```
show debugging ipv6 ospf
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show debugging ipv6 ospf

OSPFv3 debugging status:
 OSPFv3 all packet debugging is on
 OSPFv3 all NFSM debugging is on
#
```

show ipv6 ospf

Use this command to display global and area information about OSPFv3.

Command Syntax

```
show ipv6 ospf (WORD|)
```

Parameters

WORD	Tag value to use as a “match” value for controlling redistribution via route maps.

Command Mode

Privileged Exec mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
OcNOS#show ipv6 ospf
 Routing Process "OSPFv3 (null)" with ID 1.2.3.4
 Process uptime is 18 hours 12 minutes
 SPF schedule delay initial 0.500 secs
 SPF schedule delay min 0.500 secs
 SPF schedule delay max 50.0 secs
 Minimum LSA interval 5 secs, Minimum LSA arrival 1 secs
 Number of incoming current DD exchange neighbors 0/5
 Number of outgoing current DD exchange neighbors 0/5
 Number of external LSA 0. Checksum Sum 0x0000
 Number of LSA originated 6
 Number of LSA received 142
 Number of areas in this router is 1
    Area BACKBONE(0)
        Number of interfaces in this area is 2(2)
        SPF algorithm executed 12 times
        Number of LSA 6.  Checksum Sum 0x317BC
        Number of Unknown LSA 0
 Dste Staus: Disabled
```

Table 3-33 explains the fields for each ospf entry.

Table 3-33: show ipv6 ospf output details

Field	Description
Routing Process with ID	OSPFv3 process identifier and router identifier.
Process uptime is	OSPFv3 process is running time.
SPF schedule delay initial	Initial SPF schedule delay.
SPF schedule delay min	Minimum delay between receiving a change to SPF calculation.

Table 3-33: show ipv6 ospf output details

Field	Description
SPF schedule delay max	Maximum delay between receiving a change to SPF calculation.
Minimum LSA interval	Minimum LSA refresh interval.
Minimum LSA arrival	Minimum time between reception of new LSAs during flooding.
Number of incoming current DD exchange neighbors	Incoming neighbor Database Descriptors and maximum concurrent DDs. Shows the count of DD that is getting received at the moment of executing the show CLI and the maximum concurrent DDs configured.
Number of outgoing current DD exchange neighbors	Outgoing neighbor Database Descriptors and maximum concurrent DDs. Shows the count of DD that is getting sent at the moment of executing the show CLI and the maximum concurrent DDs configured.
Number of external LSA	Number of AS external LSAs and checksum.
Number of opaque AS LSA	Number of AS opaque LSAs and checksum.
Number of LSA originated	LSAs originated by the OSPFv3 instance.
Number of LSA received	LSAs received by the OSPFv3 instance.
Number of areas in this router is	Number of areas attached to this router.
Area	Area identifier.
(BACKBONE)	Area is a backbone.
no-summary	Area is a stub and does no import summaries.
(Inactive)	Area is not active.
Number of interfaces in this area is	Number of interfaces in this area.
Dste Status	Whether DSTE is enabled or disabled.

show ipv6 ospf database

Use this command to display information in the OSPFv3 Link State database.

Command Syntax

```
show ipv6 ospf database
show ipv6 ospf database (self-originate|max-age|adv-router A.B.C.D|)
show ipv6 ospf database (router|network|inter-prefix|inter-router|external|nssa-
  external|link|intra-prefix|te|grace)
show ipv6 ospf database (router|network|inter-prefix|inter-router|external|nssa-
  external|link|intra-prefix|te|grace) (self-originate|adv-router A.B.C.D|)
show ipv6 ospf database (router|network|inter-prefix|inter-router|external|nssa-
  external|link|intra-prefix|te|grace) A.B.C.D (self-originate|adv-router A.B.C.D|)
show ipv6 ospf WORD database
show ipv6 ospf WORD database (router|network|inter-prefix|inter-
  router|external|nssa-external|link|intra-prefix|te|grace)
show ipv6 ospf WORD database (router|network|inter-prefix|inter-
  router|external|nssa-external|link|intra-prefix|te|grace) adv-router A.B.C.D
```

Parameters

`self-originate`	Self-originated link states
`max-age`	LSAs in MaxAge list
`adv-router`	Advertising router for Type 8 Link LSAs (Link State Advertisements):
`A.B.C.D`	Router ID of the advertising router.
`router`	Router LSAs.
`network`	Network LSAs.
`inter-prefix`	Inter-Area-Prefix LSAs.
`inter-router`	Inter-Area-Router LSAs.
`external`	AS external LSAs.
`nssa-external`	NSSA LSAs.
`link`	Link LSAs.
`intra-prefix`	Intra-Area-Prefix LSAs (Type 9) with prefixes for stub and transit networks
`te`	Intra-area TE LSAs.
`grace`	Grace LSAs.
`A.B.C.D`	Link state ID as an IP address.
`WORD`	Tag value to use as a “match” value for controlling redistribution via route maps.

Command Mode

Privileged Exec mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example: adv-router Parameter

This example shows using the `adv-router` parameter:

```
#show ipv6 ospf database link adv-router 10.70.0.58
            OSPFv3 Router with ID (10.70.0.58) (Process 100)
                Link-LSA (Interface eth1)
  LS age: 492
  LS Type: Link-LSA
  Link State ID: 0.0.0.3
  Advertising Router: 10.70.0.58
  LS Seq Number: 0x80000001
  Checksum: 0xC2D6
  Length: 68
  Priority: 1
  Options: 0x000013 (-|R|-|-|E|V6)
  Link-Local Address: fe80::204:75ff:feaa:fedb
  Number of Prefixes: 2
    Prefix: 5f00:1:2:10::/64
    Prefix Options: 0 (-|-|-|-)
```

Header

```
OSPFv3 Router with ID (10.70.0.58) (Process 100)
Link-LSA (Interface eth1)
```

The router ID and OSPFv3 process tag of the local router.

Interface name of the router associated with this Link-LSA.

OSPFv3 Database Fields

Table 3-34 explains the fields for each database entry.

Table 3-34: OSPFv3 database fields

Field	Description
LS age	The length of time in seconds since the LSA was originated.
LS Type	The type of LSA
Link State ID	Interface identifier of the originating router.
Advertising router	The Router ID of the router advertising this LSA. On a transit network, this is always the Designated Router ID.
LS Seq Number	Sequence number of an LSA.
Checksum	LSA header checksum (excluding the LS age field).
Length	The length in bytes of the LSA (including the 20-byte header).
Priority	The router priority of the interface attaching the originating router of the link.

Table 3-34: OSPFv3 database fields

Field	Description
Options	Bits in network LSAs that originate on the link: DC-bit: Whether the router supports OSPF over Demand Circuits. R-bit: Whether the router is active. If this bit is clear, routes which transit the advertising node cannot be computed. N-bit: How the router handles Type 7 LSAs. MC-bit: Whether IP multicast packets are forwarded. E-bit: Whether AS-External-LSAs are flooded. This bit is set in all AS External LSAs and in all LSAs. originated in the backbone and non-stub areas. V6-bit: Whether to include the router/link in routing calculations.
Link-Local Address	The originating router's link-local interface address.
Number of Prefixes	The number of IPv6 prefixes associated to the link: Prefix: The global IPv6 prefix associated to this link. Prefix Options: Each prefix is advertised along with an 8-bit capabilities field.They serve as input for routing calculations allowing, for example, some prefixes to be ignored or marked as not re-advertisable.
Referenced LS Type	Identifies the Router-LSA or Network-LSA with which the IPv6 prefixes are associated: Type 0x2001: prefixes associated with Router-LSA Type 0x2002: prefixes associated with Network-LSA
Referenced Link State ID	Referenced LS Type 0x2001: this field is 0 Referenced LS Type 0x2002: the interface ID of the link's Designated Router.
Referenced Advertising Router	Referenced LS Type 0x2001: ID of the originating router. Referenced LS Type 0x2002: ID of the Designated Router Prefix: Referenced LS Type 0x2001: global IPv6 prefix associated with the router Referenced LS Type 0x2002: global IPv6 prefix associated with the transit link Prefix Options: Bits in network LSAs that originate on the link: DC: How the router handles demand circuits R: Whether the router is active. If this bit is clear, routes which transit the advertising node cannot be computed. N: How the router handles Type 7 LSAs MC: Whether IP multicast packets are forwarded E: Whether AS-External-LSAs are flooded V6: Whether to include the router/link in routing calculations Metric: The cost of this prefix.

Example: intra-prefix and adv-router Parameters

This example shows using the `adv-router` and `intra-prefix` parameters.

Note: The same information for OSPFv2 can be viewed in type 1 router LSAs and type 2 network LSAs. However, in OSPFv3 all addressing information has been removed from router LSAs and network LSAs, leading to the introduction of the Intra-Area-Prefix LSA. In a transit network, the Intra-Area-Prefix-LSA serves the same purpose as a network LSA and on a point-to-point or point-to-multipoint network serves the same purpose as a router LSA.

```
#show ipv6 ospf database intra-prefix adv-router 10.70.0.58
           OSPFv3 Router with ID (10.70.0.58) (Process 100)
```

```
               Intra-Area-Prefix-LSA (Area 0.0.0.0)
  LS age: 1435
  LS Type: Intra-Area-Prefix-LSA
  Link State ID: 0.0.0.2
  Advertising Router: 10.70.0.58
  LS Seq Number: 0x80000001
  Checksum: 0x1B4E
  Length: 56
  Number of Prefixes: 2
  Referenced LS Type: 0x2002
  Referenced Link State ID: 0.0.0.3
  Referenced Advertising Router: 10.70.0.58
    Prefix: 5f00:1:2:10::/64
    Prefix Options: 0 (-|-|-|-)
    Metric: 0
    Prefix: 6f00:1:2:10::/64
    Prefix Options: 0 (-|-|-|-)
    Metric: 0
```

Header

```
OSPFv3 Router with ID (10.70.0.58) (Process 100)
Intra-Area-Prefix-LSA (Area 0.0.0.0)
```

- The router ID and OSPFv3 process tag for the router.
- Intra-Area-Prefix-LSA has area flooding scope. This LSA belongs to Area 0.0.0.0.

show ipv6 ospf interface

Use this command to display OSPFv3 interface information.

Command Syntax

```
show ipv6 ospf interface
show ipv6 ospf interface IFNAME
```

Parameters

IFNAME	The name of the interface.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Usage

This is a sample output from the `show ipv6 ospf interface` command displaying the OSPFv3 interface information:

```
#show ipv6 ospf interface
eth0 is up, line protocol is up
  Interface ID 3, Instance ID 0, Area 0.0.0.0
  IPv6 Link-Local Address fe80::248:54ff:fec0:f32d/10
 Router ID 1.2.3.4, Network Type BROADCAST, Cost: 10
 Transmit Delay is 1 sec, State Backup, Priority 1
 Designated Router (ID) 5.6.7.8
Interface Address fe80::203:47ff:fe4c:776e
 Backup Designated Router (ID) 1.2.3.4
Interface Address fe80::248:54ff:fec0:f32d
 Timer interval configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:01
 Neighbor Count is 1, Adjacent neighbor count is 1
```

If Hello Suppression is enabled

```
RTR_B#show ipv6 ospf interface
eth1 is up, line protocol is up
  Interface ID 3
  IPv6 Prefixes
    fe80::5054:ff:fef3:f166/64 (Link-Local Address)
    2001::2/64
  OSPFv3 Process (1), Area 0.0.0.0, Instance ID 66
    Router ID 2.2.2.2, Network Type P2MP-NBMA, Cost: 1, TE Metric: 1
    Reduce LSA Flooding
    Transmit Delay is 1 sec, State Point-To-Point, Priority 1
    Timer interval configured, Hello 30, Dead 120, Wait 120, Retransmit 5
      Hello due in 00:00:32
    Neighbor Count is 1, Adjacent neighbor count is 1
    Suppress hello for 1 neighbor(s)
```

```
Hello received 2 sent 3, DD received 4 sent 6
LS-Req received 1 sent 1, LS-Upd received 7 sent 4
LS-Ack received 0 sent 3, Discarded 0
```

Table 3-35 explains the fields for each ospf interface entry.

Table 3-35: show ipv6 ospf interface output details

Field	Description
Interface Type and whether it is up or down.	Status of the interface type.
Line protocol	Status of the line protocol.
Interface ID	Interface for which information is displayed.
Instance ID	For running multiple instances of OSPFv3 on the router
Area	Area ID in A.B.C.D form
IPv6 Link-Local Address	link-local address is an IPv6 unicast address – cannot communicate to link-local addresses that are outside the directly connected network. In IPv6 (X:X::X:X/M) form.
Router ID	As stated – In A.B.C.D form.
Network Type	One of the following: • Ethernet is Broadcast • Serial p2p non-broadcast • NBMA – Non-Broadcast MultiAccess (NBMA) media
cost	The cost of sending packets over this interface – range is 1 to 65535.
Transmit Delay	The delay, in seconds, between link-state transmits. This value must be the same for all nodes on the network. The range is 1 to 65535. The default is 1.
Priority	OSPFv3 router priority for the interface. The range is 0 to 255. A router with priority 0 can never become the designated router, the default is 1.
Designated Router (ID)	The ID number of the Designated Router (DR).
Interface Address	The IPV6 address of this device.
Backup Designated Router	The ID number or the Backup Designated Router (BDR).
Interface Address	The IPv6 address of the Backup Designated Router.
Timer interval configured	The timer values of the following instances: Hello, Dead, Wait, Retransmit.
Hello due in	The countdown for receiving the next Hello packet.
Neighbor Count is	Number of neighbor count.
Adjacent neighbor count is	Number of adjacent neighbor count.

show ipv6 ospf neighbor

Use this command to display information about an OSPFv3 neighbor.

Command Syntax

```
show ipv6 ospf neighbor
show ipv6 ospf WORD neighbor
show ipv6 ospf neighbor INTERFACE
show ipv6 ospf WORD neighbor INTERFACE
show ipv6 ospf neighbor INTERFACE detail
show ipv6 ospf WORD neighbor INTERFACE detail
show ipv6 ospf neighbor detail
show ipv6 ospf WORD neighbor detail
show ipv6 ospf neighbor A.B.C.D
show ipv6 ospf WORD neighbor A.B.C.D
```

Parameters

WORD	Tag value to use as a “match” value for controlling redistribution via route maps.
INTERFACE	Display the name of the Interface
A.B.C.D	Neighbor IP address.
detail	Details of neighbors

Command Mode

Privileged Exec mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output from the show ipv6 ospf neighbor command displaying information about the OSPFv3 neighbor.

```
#show ipv6 ospf neighbor
OSPFv3 Process (*null*)
Neighbor ID Pri State Dead Time Interface Instance ID
5.6.7.8 1 Full/DR 00:00:38 eth0 0
```

If Hello Suppression is enabled

```
RTR_B#
RTR_B#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID     Pri   State           Dead Time   Interface  Instance ID
1.1.1.1           1   Full/ -         inactive    eth1       0
4.4.4.4           1   Full/DR         00:00:40    eth2       0
4.4.4.4           1   Full/ -         inactive    VLINK1     0
```

```
RTR_B#

RTR_B#
RTR_B#show ipv6 ospf neighbor detail
 Neighbor 1.1.1.1, interface address fe80::5054:ff:feb3:d3bc
    In the area 0.0.0.0 via interface eth1
    Neighbor priority is 1, State is Full, 7 state changes
    Hello is suppressed
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x000133 (AF|*|*|DC|R|-|-|E|V6)
    Dead timer due in inactive
    Database Summary List 0
    Link State Request List 0
```

Table 3-36 explains the fields for each ospf neighbor entry.

Table 3-36: show ipv6 ospf neighbor output details

Field	Description
Neighbor	Router ID of the neighbor router.
interface address	IPv6 address of the neighbor's interface.
In the area	The neighbor router's area ID.
via interface	Neighbor router's interface name.
Neighbor Priority is	OSPFv3 router priority for the interface. The range is 0 to 255. A router with priority 0 can never become the designated router, the default is 1.
State	The Link State Address (LSA) of the neighbor, and there has been 7 state changes, and sending Hello packets is suppressed.
DR	Designated Router (DR) ID
BDR	Backup Designated Router (BDR) ID
Options is	The hexadecimal representation of the seven bits in the Options Field of Hello packets (see RFC 5340): • AF-bit – Address Family bit. • V6-bit – If this bit is clear, the router/link should be excluded from IPv6 routing calculations. • E-bit – This bit describes the way AS-external-LSAs are flooded. • N-bit – This bit indicates whether or not the router is attached to an NSSA. • R-bit – This bit (the `Router' bit) indicates whether the originator is an active router. If the router bit is clear, then routes that transit the advertising node cannot be computed. Clearing the router bit is appropriate for a multi-homed host that wants to participate in routing, but does not want to forward non-locally addressed packets. • DC-bit – This bit describes the router's handling of demand circuits, as specified in [DEMAND]. • *-bit – These bits are reserved for migration of OSPFv2 protocol extensions.

Table 3-36: show ipv6 ospf neighbor output details (Continued)

Field	Description
Dead timer due in	The countdown timer for marking neighbor connections dead. In this example, the Deat Timer has been deactivated.
Database Summary List	Describes routes to IPv6 address prefixes that belong to other areas.
Link State Request List	Sent or received when Link-State Request packets finds that parts of the Link State Database are out of date.
Timer interval configured	The set values for the following packet types: Hello, Dead, Wait, Retransmit.
Neighbor Count	The number of known neighbors.
Adjacent neighbor count	The number of directly adjacent neighbors.

show ipv6 ospf route

Use this command to display the IPv6 routing table for OSPFv3.

The routes can be displayed in two ways:

- Each routing entry in a single-line
- Each routing entry in multiple lines

By default, the routing table is displayed in the multi-line format. For a single line display, give the ipv6 ospf display route single-line command.

Command Syntax

```
show ipv6 ospf route
show ipv6 ospf WORD route
```

Parameters

WORD	Tag value to use as a “match” value for controlling redistribution via route maps.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is sample output in single-line format:

```
#show ipv6 ospf route
Destination Metric Next-hop
3ffe:1:1::/48 10 directly connected, eth0
3ffe:2:1::/48 10 directly connected, eth0
3ffe:2:2::/48 10 directly connected, eth0
3ffe:3:1::/48 10 directly connected, eth0
3ffe:3:2::/48 10 directly connected, eth0
3ffe:3:3::/48 10 directly connected, eth0
E2 3ffe:100:1::1/128 10/20 via fe80::203:47ff:fe4c:776e, eth0
E2 3ffe:100:2::1/128 10/20 via fe80::203:47ff:fe4c:776e, eth0
E2 3ffe:100:3::1/128 10/20 via fe80::203:47ff:fe4c:776e, eth0
IA 3ffe:101:1::/48 20 via fe80::203:47ff:fe4c:776e, eth0
IA 3ffe:101:2::/48 20 via fe80::203:47ff:fe4c:776e, eth0
IA 3ffe:101:3::/48 20 via fe80::203:47ff:fe4c:776e, eth0
```

The following is sample output in multi-line format:

```
#show ipv6 ospf route
Destination Metric
Next-hop Interface
3ffe:1:1::/48 10
-- eth0
3ffe:2:1::/48 10
-- eth0
3ffe:2:2::/48 10
```

```
-- eth0
3ffe:3:1::/48 10
-- eth0
3ffe:3:2::/48 10
-- eth0
3ffe:3:3::/48 10
-- eth0
E2 3ffe:100:1::1/128 10/20
fe80::203:47ff:fe4c:776e eth0
```

Table 3-37 explains the fields for each ospf route entry.

Table 3-37: show ipv6 ospf route output details

Field	Description
IP address	IP address of the remote network.
Metric	For OSPF the metric is cost, which indicates the best quality path to use to forward packets.
Next hop router IP address	This route is available through the next hop router located at this IP address. This identifies exactly where packets go when they match this route.
Outgoing interface name	Interface used to get to the next-hop address for this route.

show ipv6 route fast-reroute

Use this command to display loop-free alternate routes with alternate next hops.

Command Syntax

```
show ipv6 route fast-reroute
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Usage

```
#show ipv6 route fast-reroute
```

show ipv6 ospfv3 topology

Use this command to display information about OSPFv3 topology for each area.

Command Syntax

```
show ipv6 ospfv3 topology
show ipv6 ospfv3 WORD topology
show ipv6 ospfv3 topology area (A.B.C.D|<0-4294967295>)
show ipv6 ospfv3 WORD topology area (A.B.C.D|<0-4294967295>)
```

Parameters

WORD	Tag value to use as a “match” value for controlling redistribution via route maps.
area	OSPFv3 area ID
A.B.C.D	OSPFv3 Area ID in IPv4 address format.
<0-4294967295>	OSPFv3 Area ID as a decimal value.

Command Mode

Privileged Exec mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ipv6 ospfv3 topology
OSPFv3 paths to Area (0.0.0.0) routers
Router ID Bits Metric Next-Hop Interface
1.2.3.4 --
5.6.7.8 E 10 5.6.7.8 eth0
```

Example

```
#show ipv6 ospfv3 topology

OSPFv3 paths to Area (0.0.0.0) routers
Router ID        Bits  Metric     Next-Hop            Interface
1.2.3.4                                               --
5.6.7.8          E     10         5.6.7.8             eth0
```

Table 3-38 explains the fields for each ospfv3 topology entry.

Table 3-38: show ipv6 ospfv3 topology output details

Field	Description
OSPFv3 path to Area	Area ID in IPv4 format.
Router ID	ID in IPv4 format,

Table 3-38: show ipv6 ospfv3 topology output details

Field	Description
Bits	Bits appended to packets: • V-bit Indicates whether the advertising router is an endpoint of a virtual link. • E-bit Indicates whether the advertising router is an Autonomous System Border Router (ASBR). • B-bit Indicates whether the advertising router is an Area Border Router (ABR). • W-bit When set, the router is a wild-card multicast receiver.
Metric	The value of ospfv3 metric.
Next-Hop	The next-hop identifier.
Interface	The interface name through which the virtual link extends.

show ipv6 ospf virtual-links

Use this command to display information about OSPFv3 virtual-links.

Command Syntax

```
show ipv6 ospf virtual-links
show ipv6 ospf WORD virtual-links
```

Parameters

WORD	Tag value to use as a “match” value for controlling redistribution via route maps.

Command Mode

Privileged Exec mode and Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ipv6 ospf virtual-links
Virtual Link VLINK1 to router 5.6.7.8 is up
Transit area 0.0.0.1 via interface eth0, instance ID 0
Local address 3ffe:1234:1::1/128
Remote address 3ffe:5678:3::1/128
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 Up
```

If Hello Suppression is enabled

```
RTR_B#show ipv6 ospf virtual-links
Virtual Link VLINK1 to router 4.4.4.4 is up
  Transit area 0.0.0.1 via interface eth2, instance ID 0
  Hello suppression Enabled
  DoNotAge LSA allowed
  Local address 2002::1/128
  Remote address 2002::2/128
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in inactive
    Adajcency state Full
RTR_B#
RTR_B#
```

Table 3-39 explains the fields for each ospf virtual-links entry.

Table 3-39: show ipv6 ospf virtual-links output details

Field	Description
Virtual Link	Virtual link name, the router ID to which it is connected, and the state of the link.
Transit area	Transit area ID, the interface it uses, and its instance ID – an Instance ID should default to 0. It is only necessary to assign a value other than 0 on those links that will contain multiple separate communities of OSPF routers.
Local address	The local IPV6 address and subnet mask.
Remote address	The remote IPv6 address and subnet mask.
Transmit Delay	The delay, in seconds, between link-state transmits. This value must be the same for all nodes on the network. The range is 1 to 65535. The default is 1. The state is point-to-point.
Timer intervals configured	The configured values in seconds of the following timers: Hello, Dead, Wait, Retransmit.
Hello due in	A countdown timer that indicates when the next Hello packet should arrive.
Adjacency State	Whether the adjacency state is either up or down.

© 2025 IP Infusion Inc. Proprietary

show ipv6 vrf

Use this command to list information about VRFs.

Command Syntax

```
show ipv6 vrf (WORD|)
```

Parameter

WORD	VPN Routing/Forwarding instance name.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following is a sample output of the `show ipv6 vrf` command displaying VRF information:

```
#show ipv6 vrf
Name              Interfaces

qa                eth0
you               eth1
myVRF             eth2
```

Table 3-40 explains the fields.

Table 3-40: show ipv6 vrf output details

Field	Description
Name	Name of the interface.
Interfaces	Type of an interface.

shutdown

Use this command to temporarily shut down a protocol in the least disruptive manner and to notify its neighbors that it is going away.

Use the `no` parameter of this command, not to temporarily shut a protocol.

Command Syntax

```
shutdown
no shutdown
```

Parameter

None

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#shutdown

#configure terminal
(config)#router ipv6 ospf
(config-router)#no shutdown
```

snmp restart ospf6

Use this command restart SNMP in OSPFv3

Command Syntax

```
snmp restart ospf6
```

Parameter

None

Default

By default, SNMP restart is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart ospf6
```

summary-address

Use this command to summarize or suppress external routes with the specified address range.

An address range is a pairing of a starting address and a mask that is almost the same as IP network number. For example:

- If the specified IPV6 address range is 2020:100:100:2000::/53, it matches 2020:100:100:2222::/64, 2020:100:100:2666::/64 and so on.
- If the specified IPV4 address range is 192.168.0.0/255.255.240.0, it matches 192.168.1.0/24, 192.168.4.0/22, 192.168.8.128/25 and so on.

Redistributing routes from other protocols into OSPF requires the router to advertise each route individually in an external LSA. Use this command to advertise one summary route for all redistributed routes covered by a specified network address and mask. This minimizes the size of the OSPF link state database.

Use the `no` form this command to remove summary addresses.

Command Syntax

```
summary-address X:X::X:X/M (not-advertise|(all-tag (<0-4294967295> ))| )
  (translate-tag (<0-4294967295>) | )
summary-address A.B.C.D/M (not-advertise|tag <0-4294967295>|)
no summary-address A.B.C.D/M
no summary-address X:X::X:X/M (not-advertise|all-tag|translate-tag)
no summary-address A.B.C.D/M (not-advertise|tag)
```

Parameters

`X:X::X:X/M`	The range of addresses given as IPv6 starting address and a mask.
`A.B.C.D/M`	The range of addresses given as IPv4 starting address and a mask.
`not-advertise`	Suppress routes that match the range.
`tag`	Tag value to use as a “match” value for controlling redistribution via route maps.
`<0-4294967295>`	Set a tag value. The default is 0.
`all-tag`	Set tag for all summarized type-5, translated type5 and type-7 LSA.
`translate-tag`	Set tag only for summarized translated type-5 LSA.

Default

By default, summary-address value is 0

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example uses the `summary-address` command to aggregate external LSAs that match the network 172.16.0.0/24 and assign a tag value of 3.

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#summary-address 2020:100:100:2000::/53 all-tag 3
```

timers spf exp

Use this command to set the Shortest-Path First (SPF) best-path schedule minimum and maximum delay between receiving a change to SPF calculation in milliseconds.

Use `no` parameter of this command to unset the SPF best-path schedule.

Command Syntax

```
timers spf exp <0-2147483647> <0-2147483647>
no timers spf exp
```

Parameters

`<0-2147483647>`

The minimum delay in milliseconds between receiving a change to SPF calculation.

`<0-2147483647>`

The maximum delay in milliseconds between receiving a change to SPF calculation.

Default

Default minimum delay: 500 milliseconds

Default maximum delay: 50000 milliseconds (50 seconds)

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 ospf
(config-router)#timers spf exp 300 300
```

CHAPTER 4 OSPFv3 Graceful Restart Commands

This chapter describes the OSPFv3 graceful restart commands.

- capability restart
- ipv6 ospf restart grace-period
- ipv6 ospf restart helper
- ipv6 ospf restart planned-only
- restart ipv6 ospf graceful

capability restart

Use this command to enable OSPFv3 graceful restart capability. If a router is not restart-enabled, it cannot enter graceful restart mode and act as a helper.

Use the `no` parameter with this command to disable the feature.

Command Syntax

```
capability restart graceful
```

```
no capability restart
```

Parameter

None

Default

By default, capability restart graceful is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf 100
(config-router)#capability restart graceful

(config)#router ipv6 ospf 100
(config-router)#no capability restart
```

ipv6 ospf restart grace-period

Use this command to enable the graceful restart feature and set the grace period for restarting the router.

If graceful restart is enabled, NSM is notified about the grace period. If the OSPF daemon unexpectedly shuts down, NSM sends this value to the OSPF daemon when it comes up again which uses this value to end the graceful state.

Use the `no` parameter with this command to revert to the default grace period.

Command Syntax

```
ipv6 ospf restart grace-period <2-1800>
no ipv6 ospf restart grace-period
```

Parameters

`<2-1800>`	Grace period in seconds.

Default

By default, grace period is 120 seconds

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 ospf restart grace-period 250
```

ipv6 ospf restart helper

Use this command to configure the helper behavior for graceful restart.

Use the `no` parameter with this command to revert to the default.

Command Syntax

```
ipv6 ospf restart helper {max-grace-period <2-1800>}
ipv6 ospf restart helper never (router-id A.B.C.D|)
no ipv6 ospf restart helper
no ipv6 ospf restart helper never
no ipv6 ospf restart helper {max-grace-period|never router-id (A.B.C.D|all)}
```

Parameters

`<2-1800>`	Help only if received grace-period is less than this value.
`A.B.C.D`	Router ID of neighbor to never to act as helper.
`never`	Prevent the neighbor from entering helper mode.
`max-grace-period`	Help only if received grace-period is less than this value.
`router-id`	Router of neighbor to never to act as helper.

Default

By default, router behave as helper. To disable it as helper, `ospf restart helper never` command should be configured. `ospf restart helper max-grace-period` – Max-grace-period to function as helper. If not configured, value will be the grace-period in restarting node.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 ospf restart helper never router-id 1.1.1.1

#configure terminal
(config)#no ipv6 ospf restart helper never
```

ipv6 ospf restart planned-only

Use this command to configure the OSPFv3 GR support for only planned restart.

Use the `no` parameter with this command to revert to the default.

Note: OSPFv3 supports only planned restart. Hence this command configuration and un-configuration does not have any functionality impact

Command Syntax

```
ipv6 ospf restart planned-only
```

```
no ipv6 ospf restart planned-only
```

Parameters

None

Default

By default, OSPFv3 supports planned restart only.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#ipv6 ospf restart planned-only

#configure terminal
(config)#no ipv6 ospf restart planned-only
```

restart ipv6 ospf graceful

Use this command to restart OSPFv3 gracefully.

After this command is executed, the router immediately shuts down. NSM is notified that OSPF has shut down gracefully. NSM preserves routes installed by OSPF until the grace period expires.

Command Syntax

```
restart ipv6 ospf graceful (grace-period <1-1800>|)
```

Parameters

<1-1800>	Grace period in seconds.

Default

By default, restart ipv6 ospf graceful is disabled. Default value of grace-period is 120 seconds.

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#restart ipv6 ospf graceful grace-period 200

#restart ipv6 ospf graceful
% Warning : ipv6 OSPF process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
```

CHAPTER 5 OSPF VPN Commands

This chapter provides an alphabetized reference of the OSPF VPN commands. It includes the following commands:

- capability vrf-lite
- router ospf vrf
- domain-id

capability vrf-lite

Use this command to enable the `vrf-lite` capability for an OSPF instance.

Use the `no` parameter with this command to disable the same for an OSPF instance.

Command Syntax

```
capability vrf-lite
no capability vrf-lite
```

Parameters

None

Default

By default, VRF lite capability for an OSPF instance is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ospf 100
(config-router)#capability vrf-lite
(config)#router ospf 100
(config-router)#no capability vrf-lite
```

router ospf vrf

Use this command to specify a VRF instance in OSPF. To use this command, you must first create a VRF Name in the NSM using the `ip vrf` command. Associate the same name with the OSPF instance using this command.

Command Syntax

```
router ospf <1-65535> WORD
```

Parameters

`<1-65535>`	Routing process ID; should be unique for each routing process.
`WORD`	Name of the VRF to associate with this OSPF instance.

Default

By default, router ospf vrf is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100 myVRF
(config-router)#
```

domain-id

Use this command to specify the domain ID for a OSPF bound to VRF.

The routes sent from OSPF to the VPN cloud are sent along with the domain ID. In this way, the domain ID acts as an identification for the route received from each OSPF domain.

Use the `no` form of this command to remove a domain ID.

Command Syntax

```
domain-id ((A.B.C.D (secondary|)) | (type (type-as|type-as4|type-back-comp) value HEX_DATA (secondary|)))
no domain-id ((A.B.C.D (secondary|)) | (type (type-as|type-as4|type-back-comp) value HEX_DATA (secondary|)))
```

Parameters

`A.B.C.D`	Domain ID in IP address format.
`secondary`	Domain ID is secondary. If not specified the domain ID is primary.
`type`	Domain type:
`type-as`	AS format. Hexadecimal value is 0x0005.
`type-as4`	AS4 format. Hexadecimal value is 0x0205.
`type-back-comp`	Used for backward compatibility. Hexadecimal value is 0x8000.
`value`	Domain ID.
`HEX_DATA`	Domain ID in hexadecimal.
`secondary`	Domain ID is secondary. If not specified the domain ID is primary

Default

No domain ID is defined.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example shows configuring a primary domain ID in IP address format.

```
#configure terminal
(config)#router ospf 100 vrf
(config-router)#domain-id 12.12.12.12
```

The following example shows configuring a secondary domain ID in IP address format.

```
#configure terminal
(config)#router ospf 100 vrf
```

```
(config-router)#domain-id 13.13.13.13 secondary
```

The following example shows configuring a primary domain ID in AS type format.

```
#configure terminal
(config)#router ospf 100 vrf
(config-router)#domain-id type type-as value 123456abcdef
```

Routing Information Protocol Configuration

CHAPTER 1 RIP

This chapter contains basic Router Information Protocol (RIP) configuration examples.

Enable RIP

This example shows the minimum configuration required to enable RIP on an interface. R1 and R2 are two routers connecting to network 10.10.11.0/24. R1 and R2 are also connected to networks 10.10.10.0/24 and 10.10.12.0/24, respectively. To enable RIP, first define the RIP routing process, then associate a network with the routing process.

Topology

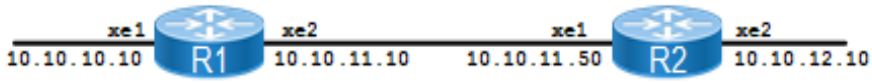

Figure 1-62: Enable RIP Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.10.0/24` `(config-router)#network 10.10.11.0/24`	Associate networks with the RIP process.
`(config-router)#exit`	Exit router mode and return to configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

R2

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.11.0/24` `(config-router)#network 10.10.12.0/24`	Associate networks with the RIP process.
`(config-router)#exit`	Exit router mode and return to configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

Validation

show ip rip, show running-config, show ip protocols rip, show ip rip interface, show ip route

R1

```
#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
```

```
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                           1                  xe1
Rc 10.10.11.0/24                           1                  xe2
R  10.10.12.0/24      10.10.11.50          2 10.10.11.50      xe2    02:32

#show running-config rip
!
router rip
 network 10.10.10.0/24
 network 10.10.11.0/24
!

#show ip protocols rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 2 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing:
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    xe48             2     2
    ce49             2     2
  Routing for Networks:
    10.10.10.0/24
    10.10.11.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.11.50           120  00:00:06               0           0
  Number of routes (including connected): 3
  Distance: (default is 120)

#show ip rip interface
lo is up, line protocol is up
  RIP is not enabled on this interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
xe1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.10/24
xe2 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.11.10/24
```

```
xe3 is up, line protocol is up
  RIP is not enabled on this interface
...

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, xe1, 00:08:01
C            10.10.11.0/24 is directly connected, xe2, 00:07:34
R            10.10.12.0/24 [120/2] via 10.10.11.50, xe2, 00:05:10
C            127.0.0.0/8 is directly connected, lo, 4d18h40m
C            192.168.0.2/32 is directly connected, lo, 4d13h46m

Gateway of last resort is not set
```

R2

```
#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
R  10.10.10.0/24      10.10.11.10           2 10.10.11.10     xe1    02:34
Rc 10.10.11.0/24                            1                 xe1
Rc 10.10.12.0/24                            1                 xe2

#show running-config rip
!
router rip
 network 10.10.11.0/24
 network 10.10.12.0/24
!

#show ip protocols rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 25 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing:
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    xe2              2     2
    ce49             2     2
  Routing for Networks:
```

```
    10.10.11.0/24
    10.10.12.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.11.10           120  00:00:13               0           0
  Number of routes (including connected): 3
  Distance: (default is 120)

#show ip rip interface
lo is up, line protocol is up
  RIP is not enabled on this interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
xe1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.11.50/24
xe2 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.12.10/24
xe3 is up, line protocol is up
  RIP is not enabled on this interface
  ...

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
R            10.10.10.0/24 [120/2] via 10.10.11.10, xe1, 00:11:08
C            10.10.11.0/24 is directly connected, xe1, 00:13:00
C            10.10.12.0/24 is directly connected, xe2, 00:12:26
C            127.0.0.0/8 is directly connected, lo, 4d18h50m
C            192.168.0.1/32 is directly connected, lo, 4d14h01m

Gateway of last resort is not set
```

Specify RIP Version

Configure a router to receive and send specific versions of packets on an interface. In this example, router R2 is configured to receive and send RIP version 1 and version 2 information on both eth1 and eth2 interfaces.

Topology

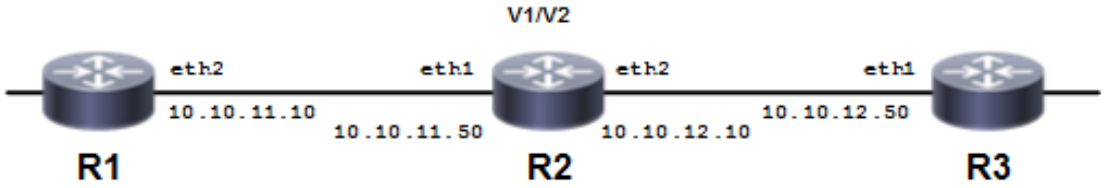

Figure 1-63: RIP Version Topology

R2

`#configure terminal`	Enter configure mode
`(config)#router rip`	Enable the RIP routing process
`(config-router)#exit`	Exit router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip rip send version 1 2`	Send RIP version 1 and version 2 packets out this interface
`(config-if)#ip rip receive version 1 2`	Receive RIP version 1 and version 2 packets from this interface
`(config-if)#exit`	Exit interface mode
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip rip send version 1 2`	Send RIP version 1 and version 2 packets out this interface
`(config-if)#ip rip receive version 1 2`	Receive RIP version 1 and version 2 packets from this interface
`(config-if)#exit`	Exit router mode and return to configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

Validation

R2

```
#sh ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.11.0/24                           1                  eth1
Rc 10.10.12.0/24                           1                  eth2
```

```
#sh running-config
!
no service password-encryption
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
spanning-tree mode provider-rstp
data-center-bridging enable
feature telnet
feature ssh
no feature tacacs+
snmp-server view all .1 included
ntp enable
sFlow disable
software-watchdog keep-alive-time 30
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ip address 192.168.0.2/32 secondary
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.4.108/24
!
interface eth1
 ip address 10.10.11.50/24
 ip rip send version 1 2
 ip rip receive version 1 2
!
interface eth2
 ip address 10.10.12.10/24
 ip rip send version 1 2
 ip rip receive version 1 2
!
router rip
 network 10.10.11.0/24
 network 10.10.12.0/24
!
line con 0
 login
line vty 0 39
 login
!
end

#show ip protocols rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 29 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
```

```
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing:
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth1             1 2   1 2
    eth2             1 2   1 2
  Routing for Networks:
    10.10.11.0/24
    10.10.12.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.11.10           120  00:00:31               0           0
    10.10.12.50           120  00:00:08               0           0
  Number of routes (including connected): 2
  Distance: (default is 120)

#show ip rip interface
svlan0.1 is down, line protocol is down
  RIP is not enabled on this interface
eth2 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIPv1 and RIPv2 packets
    Send RIPv1 and RIPv2 packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.12.10/24
eth1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIPv1 and RIPv2 packets
    Send RIPv1 and RIPv2 packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.11.50/24
eth0 is up, line protocol is up
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.11.0/24 is directly connected, eth1, 00:04:22
C            10.10.12.0/24 is directly connected, eth2, 00:10:59
C            127.0.0.0/8 is directly connected, lo, 4d19h04m
C            192.168.0.1/32 is directly connected, lo, 4d14h15m
```

Authentication with a Single Key

OcNOS RIP provides a choice of configuring authentication with a single key or with multiple keys. This example shows authenticating routing information exchange using a single key.

Topology

Routers R1 and R2 are running RIP and exchanging routing updates. To configure single-key authentication on R1, specify an interface, then define a key or password for that interface. Next, specify an authentication mode. Any receiving RIP packet on this specified interface should have the same string as the password. For an exchange of updates between R1 and R2, define the same password and authentication mode on R2.

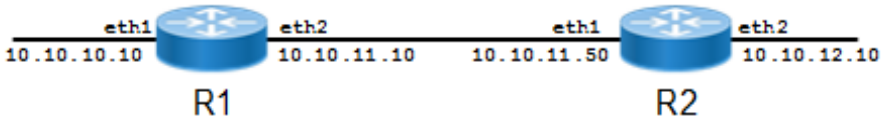

Figure 1-64: Single-key Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.10.0/24`	Associate network `10.10.10.0/24` with the RIP process.
`(config-router)#redistribute connected`	Enable redistributing from connected routes.
`(config-router)#exit`	Exit router mode.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#interface eth1`	Specify the interface (`eth1`) for authentication.
`(config-if)#ip rip authentication string ABC`	Specify the authentication string (`ABC`) for this interface.
`(config-if)#ip rip authentication mode md5`	Specify the authentication mode to be MD5.
`(config-if)#exit`	Exit router mode and return to configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

R2

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.11.0/24`	Associate network `10.10.11.0/24` with the RIP process.
`(config-router)#redistribute connected`	Enable redistributing from connected routes.
`(config-router)#exit`	Exit router mode.
`(config)#commit`	Commit the candidate configuration to the running configuration

(config)#interface eth2	Specify the interface (eth2) for authentication.
(config-if)#ip rip authentication string ABC	Specify the authentication string (ABC) on this interface.
(config-if)#ip rip authentication mode md5	Specify the authentication mode to be MD5.
(config-if)#exit	Exit router mode and return to configure mode.
(config)#commit	Commit the candidate configuration to the running configuration

Validation

show running-config, show ip rip, show ip protocol rip, show ip rip interface, show ip route

R1

```
#show running-config
!
no service password-encryption
!
hostname rtr1
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
spanning-tree mode provider-rstp
data-center-bridging enable
feature telnet
feature ssh
no feature tacacs+
snmp-server view all .1 included
ntp enable
sFlow disable
software-watchdog keep-alive-time 30
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ip address 192.168.0.1/32 secondary
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.4.92/24
!
interface eth1
 ip address 10.10.10.10/24
 ip rip authentication mode md5
 ip rip authentication string 0x5c5b790e25d29287
!
interface eth2
 ip address 10.10.11.10/24
!
router rip
```

```
 network 10.10.10.0/24
 redistribute connected
!
line con 0
 login
line vty 0 39
 login
!
end

#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                           1                  eth1
Rc 10.10.11.0/24                           1                  eth2
R  10.10.12.0/24      10.10.11.50          2 10.10.11.50      eth2   02:41
R  192.168.0.1/32     10.10.11.50          2 10.10.11.50      eth2   02:41
C  192.168.0.2/32                          1                  lo

#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 26 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth1             2     2
  Routing for Networks:
    10.10.10.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.50           120  00:00:31               0           0
  Number of routes (including connected): 6
  Distance: (default is 120)

#show ip rip interface
svlan0.1 is down, line protocol is down
  RIP is not enabled on this interface
eth2 is up, line protocol is up
  RIP is not enabled on this interface
eth1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.10/24
eth0 is up, line protocol is up
```

```
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.4.1 to network 0.0.0.0

K*      0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C       10.10.10.0/24 is directly connected, eth1
C       10.10.11.0/24 is directly connected, eth2
R       10.10.12.0/24 [120/2] via 10.10.10.50, eth1, 00:04:05
C       10.12.4.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo
C       192.168.0.1/32 is directly connected, lo
R       192.168.0.2/32 [120/2] via 10.10.10.50, eth1, 00:04:05
```

R2

```
#sh running-config
!
no service password-encryption
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
spanning-tree mode provider-rstp
data-center-bridging enable
feature telnet
feature ssh
no feature tacacs+
snmp-server view all .1 included
ntp enable
sFlow disable
software-watchdog keep-alive-time 30
!
ip pim register-rp-reachability
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ip address 192.168.0.2/32 secondary
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.4.108/24
!
interface eth1
```

```
 ip address 10.10.12.50/24
!
interface eth2
 ip address 10.10.10.50/24
 ip rip authentication mode md5
 ip rip authentication string 0x5c5b790e25d29287
!
router rip
 network 10.10.10.0/24
 redistribute connected
!
line con 0
 login
line vty 0 39
 login
!
end

#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
R  10.10.10.0/24      10.10.11.10           2 10.10.11.10     eth1   02:37
Rc 10.10.11.0/24                            1                 eth1
Rc 10.10.12.0/24                            1                 eth2
C  192.168.0.1/32                           1                 lo
R 192.168.0.2/32      10.10.11.10           2 10.10.11.10     eth1   02:37

#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 5 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth2             2     2
  Routing for Networks:
    10.10.10.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.10           120  00:00:01               0           0
  Number of routes (including connected): 6
  Distance: (default is 120)

#show ip rip interface
svlan0.1 is down, line protocol is down
  RIP is not enabled on this interface
eth2 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
```

```
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.50/24
eth1 is up, line protocol is up
  RIP is not enabled on this interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface

#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default

IP Route Table for VRF "default"
Gateway of last resort is 10.12.4.1 to network 0.0.0.0

K*      0.0.0.0/0 [0/0] via 10.12.4.1, eth0
C       10.10.10.0/24 is directly connected, eth2
R       10.10.11.0/24 [120/2] via 10.10.10.10, eth2, 00:07:36
C       10.10.12.0/24 is directly connected, eth1
C       10.12.4.0/24 is directly connected, eth0
C       127.0.0.0/8 is directly connected, lo
R       192.168.0.1/32 [120/2] via 10.10.10.10, eth2, 00:07:36
C       192.168.0.2/32 is directly connected, lo
```

Text Authentication with Multiple Keys

This example illustrates text authentication of the routing information exchange process for RIP using multiple keys.

Topology

Routers R1 and R2 are running RIP, and exchanging routing updates. To configure authentication on R1, define a key chain, specify keys in the key chain, then define the authentication string or passwords to use by the keys. Set the time period during which it is valid to receive or send the authentication key by specifying the accept and send lifetimes. After defining the key string, specify the key chain (or set of keys) that will be used for authentication on each interface, and the authentication mode to use.

R1 receives all packets that contain any key string that matches one of the key strings included in the specified key chain (within the accept lifetime) on that interface. The key ID is not considered for matching. For additional security, the accept lifetime and send lifetime are configured such that every fifth day, the key ID and key string changes. To maintain continuity, the accept lifetimes should be configured to overlap. This will accommodate different time setup on machines. However, the send lifetime is not required to overlap, and IP Infusion Inc. recommends configuring no overlapping for the send lifetime.

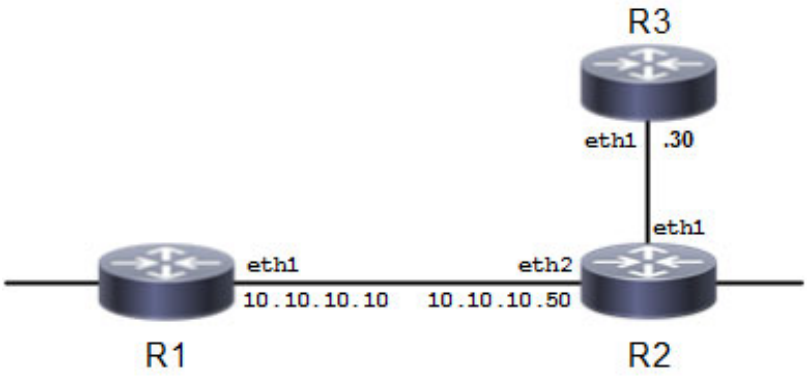

Figure 1-65: Multiple-key Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.10.0/24`	Associate network `10.10.10.0/24` with the RIP process.
`(config-router)#redistribute connected`	Enable redistributing from connected routes.
`(config-router)#exit`	Exit router mode.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#key chain SUN`	Enter Keychain management mode to add keys to the key chain `SUN`.
`(config-keychain)#key-id 10`	Add authentication key ID (`10`) to the key chain `SUN`.
`(config-keychain-key)#key-string ABC`	Specify a password (`ABC`) to use by the specified key.
`(config-keychain-key)#accept-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024`	Specify the time period during which the authentication key can be received. In this case, key string `ABC` can be received from 7 PM of Aug 27 to 11 PM Aug 31, 2024.
`(config-keychain-key)#send-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024`	Specify the time period during which the authentication key can be sent. In this case, key string `ABC` can be sent from 7 PM of Aug 27 to 11 PM Aug 31, 2024.
`(config-keychain-key)#exit`	Exit Keychain-Key mode, and return to Keychain mode.
`(config-keychain)#commit`	Commit the candidate configuration to the running configuration
`(config-keychain)#key-id 20`	Add another authentication key (`20`) to the key chain `SUN`.
`(config-keychain-key)#key-string Earth`	Specify a password (`Earth`) to use by the specified key.
`(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024`	Specify the time period during which authentication key string `Earth` can be received. In this case, key string `Earth` can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.

(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which the authentication key can be sent. In this case, key string Earth can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#commit	Commit the candidate configuration to the running configuration
(config-keychain-key)#exit	Exit Keychain-Key mode.
#configure terminal	Enter configure mode.
(config)#interface eth1	Specify interface eth1 as the interface you want to configure.
(config-if)#ip address 10.10.10.10/24	Assign the IP address to an interface eth1.
(config-if)#ip rip authentication key-chain SUN	Enable RIPv2 authentication on eth1 interface and specify the key-chain SUN to use for authentication.
(config-if)#ip rip authentication mode text	Specify text authentication mode to use for RIP packets. This step is optional, because text is the default mode.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#interface eth2	Specify interface eth2 as the interface you want to configure.
(config-if)#ip address 44.4.4.4/24	Assign the IP address to an interface eth2.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

R2

#configure terminal	Enter configure mode.
(config)#router rip	Define a RIP routing process, and enter Router mode.
(config-router)#network 10.10.10.0/24	Associate network 10.10.10.0/24 with the RIP process.
(config-router)#redistribute connected	Enable redistributing from connected routes.
(config-router)#exit	Exit router mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#key chain MOON	Enter Keychain management mode to add keys to the key chain MOON.
(config-keychain)#key-id 30	Add authentication key ID (30) to the key chain MOON.
(config-keychain-key)#key-string ABC	Specify a password (ABC) to use by the specified key.
(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which authentication key string ABC can be received. In this case, key string ABC can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which the authentication key can be sent. In this case, key string ABC can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#exit	Exit Keychain-Key mode, and return to Keychain mode.

(config-keychain)#commit	Commit the candidate configuration to the running configuration
(config-keychain)#key-id 40	Add another authentication key (40) to the key chain MOON.
(config-keychain-key)#key-string Earth	Specify a password (Earth) to use by the specified key.
(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which authentication key string Earth can be received. In this case, key string Earth can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which the authentication key can be sent. In this case, key string Earth can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#commit	Commit the candidate configuration to the running configuration
(config-keychain-key)#exit	Exit Keychain-Key mode.
#configure terminal	Enter configure mode.
(config)#interface eth2	Specify interface eth2 as the interface you want to configure.
(config-if)#ip address 10.10.10.50/24	Assign the IP address to an interface eth2.
(config-if)#ip rip authentication key-chain MOON	Enable RIPv2 authentication on the eth2 interface, and specify the key-chain MOON to use for authentication.
(config-if)#ip rip authentication mode text	Specify the authentication mode to use for RIP packets. This step is optional, because text is the default mode.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#interface eth1	Specify interface eth1 as the interface you want to configure.
(config-if)#ip address 55.5.5.5/24	Assign the IP address to an interface eth1.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

Validation

show running-config, show ip rip, show ip protocol rip, show ip rip interface, show ip route

R1

Here is the snippet configuration for R1 in the given network topology.

```
R1#show running-config
!
key chain SUN
 key-id 10
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024
```

```
 key-id 20
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ip address 192.168.0.1/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.4.92/24
!
interface eth1
 ip address 10.10.10.10/24
 ip rip authentication mode text
 ip rip authentication key-chain SUN
!
interface eth2
  ip address 44.4.4.4/24
!
router rip
 network 10.10.10.0/24
 redistribute connected
!
!
end

R1#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                         1                   eth1
C  44.4.4.0/24                           1                    eth2
R  55.5.5.0/24        10.10.10.50        2 10.10.10.50       eth1   02:29
C  192.168.0.1/32                        1                   lo

#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 4294967295 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
```

```
    eth1              2     2      SUN
  Routing for Networks:
    10.10.10.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.50           120  00:00:21               0           0
  Number of routes (including connected): 4
  Distance: (default is 120)

#show ip rip interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface
lo.management is up, line protocol is up
  RIP is not enabled on this interface
eth1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.10/24
eth2 is up, line protocol is up
  RIP is not enabled on this interface
```

```
R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth1, installed 00:58:03, last update 00:58:03 ago
C            44.4.4.0/24 is directly connected, eth2, installed 00:52:08, last update 00:52:08 ago
R            55.5.5.0/24 [120/2] via 10.10.10.50, eth1, installed 00:08:12, last update 00:08:12 ago
C            127.0.0.0/8 is directly connected, lo, installed 01:27:43, last update 01:27:43 ago
C            192.168.0.1/32 is directly connected, lo, installed 01:13:23, last update 01:13:23 ago

Gateway of last resort is not set
```

R2

Here is the snippet configuration for R2 in the given network topology.

```
R2#show running-config
!
key chain MOON
 key-id 30
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
 key-id 40
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
!
```

```
interface eth0
 ip address 10.12.4.0/24
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth1
 ip address 55.5.5.5/24
!
interface eth2
 ip address 10.10.10.50/24
 ip rip authentication mode text
 ip rip authentication key-chain MOON
!
router rip
 network 10.10.10.0/24
 redistribute connected
!
!
end

R2#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                           1                  eth2
R  44.4.4.0/24        10.10.10.10          2 10.10.10.10      eth2   02:40
C  55.5.5.0/24                             1                  eth1
R  192.168.0.1/32     10.10.10.10          2 10.10.10.10      eth2   02:40

R2#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 12 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth2             2     2      MOON
  Routing for Networks:
    10.10.10.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.10           120  00:00:25               0           0
  Number of routes (including connected): 4
```

```
      Distance: (default is 120)

    R2#show ip rip interface
    eth0 is up, line protocol is up
      RIP is not enabled on this interface
    lo is up, line protocol is up
      RIP is not enabled on this interface
    lo.management is up, line protocol is up
      RIP is not enabled on this interface
    eth1 is up, line protocol is up
      RIP is not enabled on this interface
    eth2 is up, line protocol is up
      Routing Protocol: RIP
        Receive RIP packets
        Send RIP packets
        Passive interface: Disabled
        Split horizon: Enabled with Poisoned Reversed
        IP interface address:
          10.10.10.50/24
```

```
R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            10.10.10.0/24 is directly connected, eth2, installed 00:59:06, last update 00:59:06 ago
R            44.4.4.0/24 [120/2] via 10.10.10.10, eth2, installed 00:02:26, last update 00:02:26 ago
C            55.5.5.0/24 is directly connected, eth1, installed 00:52:46, last update 00:52:46 ago
C            127.0.0.0/8 is directly connected, lo, installed 01:23:42, last update 01:23:42 ago
R            192.168.0.1/32 [120/2] via 10.10.10.10, xe25, installed 00:02:26, last update 00:02:26 ago

Gateway of last resort is not set
```

MD5 Authentication with Multiple Keys

This example illustrates the MD5 authentication of the routing information exchange process for RIP using multiple keys.

Topology

Routers R1 and R2 are running RIP, and exchanging routing updates. To configure authentication on R1, define a key chain, specify keys in the key chain, then define the authentication string or passwords to use by the keys. Then, set the time period during which it is valid to receive or send the authentication key by specifying the accept and send lifetimes. After defining the key string, specify the key chain (or the set of keys) that will be used for authentication on the interface, and the authentication mode to use. Configure R2 and R3 to have the same key ID and key string as R1 for the time that updates are to be exchanged.

In MD5 authentication, both the key ID and key string are matched for authentication. R1 will receive only packets that match both the key ID and the key string in the specified key chain (within the accept lifetime) on that interface. In the following example, R2 has the same key ID and key string as R1. For additional security, the accept lifetime and send lifetime are configured such that every fifth day, the key ID and key string changes. To maintain continuity, the accept lifetimes should be configured to overlap; however, the send lifetime should not overlap.

Figure 1-66: MD5 Multiple-key Topology

R1

`#configure terminal`	Enter configure mode.
`(config)#router rip`	Define a RIP routing process, and enter Router mode.
`(config-router)#network 10.10.10.0/24`	Associate network `10.10.10.0/24` with the RIP process.
`(config-router)#redistribute connected`	Enable redistributing from connected routes.
`(config-router)#exit`	Exit router mode.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#key chain SUN`	Enter Keychain management mode to add keys to the key chain `SUN`.
`(config-keychain)#key-id 1`	Add authentication key ID (`1`) to the key chain `SUN`.
`(config-keychain-key)#key-string ABC`	Specify a password (`ABC`) to use by the specified key.
`(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024`	Specify the time period during which authentication key string `ABC` can be received. In this case, key string `ABC` can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
`(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024`	Specify the time period during which the authentication key can be sent. In this case, key string `ABC` can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
`(config-keychain-key)#exit`	Exit Keychain-Key mode, and return to Keychain mode.
`(config-keychain)#commit`	Commit the candidate configuration to the running configuration
`(config-keychain)#key-id 2`	Add another authentication key (`2`) to the key chain `SUN`.
`(config-keychain-key)#key-string Earth`	Specify a password (`Earth`) to use by the specified key.
`(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024`	Specify the time period during which authentication key string `Earth` can be received. In this case, key string `Earth` can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
`(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024`	Specify the time period during which the authentication key can be sent. In this case, key string `Earth` can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.

(config-keychain-key)#commit	Commit the candidate configuration to the running configuration
(config-keychain-key)#exit	Exit Keychain-Key mode.
#configure terminal	Enter configure mode.
(config)#interface eth1	Specify interface eth1 as the interface you want to configure.
(config-if)#ip address 10.10.10.10/24	Assign the IP address to an interface eth1.
(config-if)#ip rip authentication key-chain SUN	Enable RIPv2 authentication on the eth1 interface, and specify the key chain SUN to use for authentication.
(config-if)#ip rip authentication mode md5	Specify MD5 authentication mode to use for RIP packets.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#interface eth2	Specify interface eth2 as the interface you want to configure.
(config-if)#ip address 44.4.4.4/24	Assign the IP address to an interface eth2.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

R2

#configure terminal	Enter configure mode.
(config)#router rip	Define a RIP routing process, and enter Router mode.
(config-router)#network 10.10.10.0/24	Associate network 10.10.10.0/24 with the RIP process.
(config-router)#redistribute connected	Enable redistributing from connected routes.
(config-router)#exit	Exit router mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#key chain MOON	Enter Keychain management mode to add keys to the key chain MOON.
(config-keychain)#key-id 1	Add authentication key ID (1) to the key chain MOON.
(config-keychain-key)#key-string ABC	Specify a password (ABC) to use by the specified key.
(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which authentication key string ABC can be received. In this case, key string ABC can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which the authentication key can be sent. In this case, key string ABC can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#exit	Exit Keychain-Key mode, and return to Keychain mode.
(config-keychain)#commit	Commit the candidate configuration to the running configuration

(config-keychain)#key-id 2	Add another authentication key (2) to the key chain MOON.
(config-keychain-key)#key-string Earth	Specify a password (Earth) to use by the specified key.
(config-keychain-key)#accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which authentication key string Earth can be received. In this case, key string Earth can be received from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024	Specify the time period during which the authentication key can be sent. In this case, key string Earth can be sent from 7 PM of Aug 02 to 11 PM Aug 31, 2024.
(config-keychain-key)#commit	Commit the candidate configuration to the running configuration
(config-keychain-key)#end	Enter Privileged Exec mode.
#configure terminal	Enter configure mode.
(config)#interface eth2	Specify interface eth2 as the interface you want to configure.
(config-if)#ip address 10.10.10.50/24	Assign the IP address to an interface eth2.
(config-if)#ip rip authentication key-chain MOON	Enable RIPv2 authentication on the eth1 interface, and specify the key chain MOON to use for authentication.
(config-if)#ip rip authentication mode md5	Specify the authentication mode to use for RIP packets.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration
(config)#interface eth1	Specify interface eth1 as the interface you want to configure.
(config-if)#ip address 55.5.5.5/24	Assign the IP address to an interface eth1.
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration

Validation

show running-config, show ip rip, show ip protocol rip, show ip rip interface

R1

Here is the snippet configuration for R1 in the given network topology.

```
R1#show running-config
!
key chain SUN
 key-id 10
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 27 2024 23:00:00 Aug 31 2024
 key-id 20
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
```

```
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
!
interface lo
 mtu 65536
 ip address 127.0.0.1/8
 ip address 192.168.0.1/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth0
 ip address 10.12.4.92/24
!
interface eth1
 ip address 10.10.10.10/24
 ip rip authentication mode md5
 ip rip authentication key-chain SUN
!
interface eth2
  ip address 44.4.4.4/24
!
router rip
 network 10.10.10.0/24
 redistribute connected
!
!
end

R1#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                         1                 eth1
C  44.4.4.0/24                           1                  eth2
R  55.5.5.0/24        10.10.10.50        2 10.10.10.50     eth1   02:29
C  192.168.0.1/32                        1                 lo

R1#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 4294967295 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth1             2     2      SUN
  Routing for Networks:
    10.10.10.0/24
```

```
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.50           120  00:00:21               0           0
  Number of routes (including connected): 4
  Distance: (default is 120)

#show ip rip interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface
lo.management is up, line protocol is up
  RIP is not enabled on this interface
eth1 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      10.10.10.10/24
eth2 is up, line protocol is up
  RIP is not enabled on this interface
```

R2

Here is the snippet configuration for R2 in the given network topology.

```
R2#show running-config
!
key chain MOON
 key-id 30
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
 key-id 40
  key-string encrypted 0xa057668002822d4f0e04131ff4996b184d8711aa527604f7
  accept-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
  send-lifetime 19:00:00 Aug 02 2024 23:00:00 Aug 31 2024
!
interface eth0
 ip address 10.12.4.0/24
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface eth1
 ip address 55.5.5.5/24
!
interface eth2
 ip address 10.10.10.50/24
 ip rip authentication mode md5
```

```
 ip rip authentication key-chain MOON
!
router rip
 network 10.10.10.0/24
 redistribute connected
!
!
end

R2#show ip rip

Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP,
       X - Default

   Network            Next Hop         Metric From            If     Time
Rc 10.10.10.0/24                           1                  eth2
R  44.4.4.0/24        10.10.10.10          2 10.10.10.10      eth2   02:40
C  55.5.5.0/24                             1                  eth1
R  192.168.0.1/32     10.10.10.10          2 10.10.10.10      eth2   02:40

R2#show ip protocol rip
RIP Database for VRF (default)
Routing Protocol is "rip"
  Sending updates every 30 seconds with +/-50%, next due in 12 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribution metric is 1
  Redistributing: connected
  Default version control: send version 2, receive version 2
    Interface        Send  Recv   Key-chain
    eth2             2     2      MOON
  Routing for Networks:
    10.10.10.0/24
  Routing Information Sources:
    Gateway          Distance  Last Update  Bad Packets  Bad Routes
    10.10.10.10           120  00:00:25               0           0
  Number of routes (including connected): 4
  Distance: (default is 120)

R2#show ip rip interface
eth0 is up, line protocol is up
  RIP is not enabled on this interface
lo is up, line protocol is up
  RIP is not enabled on this interface
lo.management is up, line protocol is up
  RIP is not enabled on this interface
eth1 is up, line protocol is up
  RIP is not enabled on this interface
eth2 is up, line protocol is up
  Routing Protocol: RIP
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
```

```
10.10.10.50/24
```

CHAPTER 2 RIPng

This chapter contains a basic RIPng configuration example.

Topology

The diagram shows the minimum configuration required to enable RIPng on an interface. R1 and R2 are two routers connected to network 3ffe:11::/64. To enable RIPng, first define the RIPng routing process, then enable RIPng on each interface.

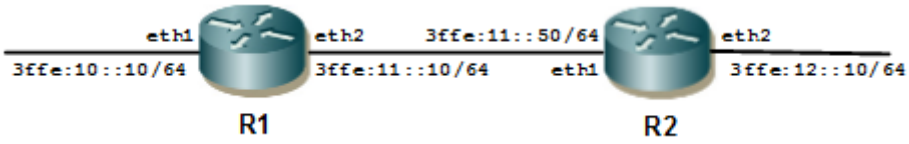

Figure 2-67: RIPng Topology

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router rip`	Enable RIPng routing on interface `eth1`.
`(config-if)#exit`	Exit interface mode, and enter Configure mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router rip`	Enable RIPng routing on interface `eth2`.
`(config-if)#exit`	Exit interface mode, and enter Configure mode.
`(config)#router ipv6 rip`	Define a RIPng routing process, and enter Router mode.

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ipv6 router rip`	Enable RIPng routing on interface `eth1`.
`(config-if)#exit`	Exit interface mode, and enter Configure mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ipv6 router rip`	Enable RIPng routing on interface `eth2`.
`(config-if)#exit`	Exit interface mode, and enter Configure mode.
`(config)#router ipv6 rip`	Define a RIPng routing process, and enter Router mode.

Validation

show ipv6 rip

Routing Information Protocol Command Reference

CHAPTER 1 Routing Information Protocol Commands

This chapter provides an alphabetized reference for each of the Routing Information Protocol (RIP) commands, which support IPv4. It includes the following commands:

- accept-lifetime
- cisco-metric-behavior
- clear ip rip route
- clear ip rip route vrf NAME
- clear ip rip statistics
- debug rip
- default-information originate
- default-metric
- distance
- distribute-list
- ip rip authentication key-chain
- ip rip authentication mode
- ip rip authentication string
- ip rip receive-packet
- ip rip receive version
- ip rip send-packet
- ip rip send version
- ip rip split-horizon
- key
- key chain
- key-string
- maximum-prefix
- neighbor
- network
- offset-list
- passive-interface
- recv-buffer-size
- redistribute
- route
- router rip
- send-lifetime
- show debugging rip
- show ip protocols rip

© 2025 IP Infusion Inc. Proprietary

- show ip rip
- show ip rip interface
- show ip rip statistics
- snmp restart rip
- timers basic
- version

accept-lifetime

Use this command to specify the time period during which the authentication key on a key chain is received as valid.

Use the `no` option with this command to disable it.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
accept-lifetime HH:MM:SS <1-31> MONTH <1993-2035> HH:MM:SS <1-31> MONTH <1993-2035>
accept-lifetime HH:MM:SS <1-31> MONTH <1993-2035> HH:MM:SS MONTH <1-31> <1993-2035>
accept-lifetime HH:MM:SS MONTH <1-31> <1993-2035> HH:MM:SS <1-31> MONTH <1993-2035>
accept-lifetime HH:MM:SS MONTH <1-31> <1993-2035> HH:MM:SS MONTH <1-31> <1993-2035>
accept-lifetime HH:MM:SS <1-31> MONTH <1993-2035> infinite
accept-lifetime HH:MM:SS MONTH <1-31> <1993-2035> infinite
accept-lifetime HH:MM:SS <1-31> MONTH <1993-2035> duration <1-2147483646>
accept-lifetime HH:MM:SS MONTH <1-31> <1993-2035> duration <1-2147483646>
no accept-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of accept-lifetime in hours, minutes and seconds.
`<1-31>`	Specify the day of the month to start.
`MONTH`	Specify the month of the year to start (the first three letters of the month, for example, `Jan.`).
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the time when accept-lifetime expires in hours, minutes and seconds.
`<1-31>`	Specify the day of the month to expire.
`MONTH`	Specify the month of the year to expire (the first three letters of the month, for example, `Jan.`).
`<1993-2035>`	Specify the year to expire.
`duration`	Specify the duration of the key in seconds <1-2147483646>.
`infinite`	Specify the end time to never expire.

Default

None

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of accept-lifetime for `key1` on the key chain named `mychain`.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 3 2004 04:04:02 Oct 6 2006

(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#no accept-lifetime
```

cisco-metric-behavior

Use this command to enable the metric update consistent with Cisco.

Use either the `no` or `disable` parameter with this command to disable this feature.

Command Syntax

```
cisco-metric-behavior (enable|disable)
no cisco-metric-behavior
```

Parameters

`enable`	Enable updating the metric consistent with Cisco.
`disable`	Disable updating the metric consistent with Cisco.

Default

By default, the Cisco metric-behavior is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to enable the metric update behavior to be consistent with Cisco in the Router mode.

```
#configure terminal
(config)#router rip
 (config-router)#cisco-metric-behavior enable
```

clear ip rip route

Use this command to clear specific data from the RIP routing tables.

Using this command with the `all` parameter, clears the RIP table of all the routes. If you do not want that your RIP network to be deleted, use the `redistribute connected` command and make the RIP network a connected route. To delete the RIP routes learned from neighbor and also keep the RIP network intact, use the `rip` (`clear ip rip route rip`) parameter with this command.

Command Syntax

```
clear ip rip route (A.B.C.D/M|rip|kernel|connected|static|ospf|isis|bgp|all)
```

Parameters

`A.B.C.D/M`	Removes entries which exactly match this destination address from RIP routing table.
`bgp`	Removes only BGP routes from the RIP routing table.
`connected`	Removes entries for connected routes from the RIP routing table.
`isis`	Removes only IS-IS routes from the RIP routing table
`kernel`	Removes kernel entries from the RIP routing table.
`ospf`	Removes only OSPF routes from the RIP routing table.
`rip`	Removes only RIP routes from the RIP routing table.
`static`	Removes static entries from the RIP routing table.
`all`	Removes the entire RIP routing table.

Default

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip rip route 10.0.0.0/8
#clear ip rip route ospf
```

clear ip rip route vrf NAME

Use this command to clear all IPv4 RIP VRF route or any specific prefix RIP VRF route of any particular VRF name.

Command Syntax

```
clear ip rip route vrf NAME (*|A.B.C.D/M)
```

Parameters

A.B.C.D/M	Removes entries with the prefix specified.
*	Removes all routes
NAME	VPN Routing or Forwarding instance name

Default

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip rip route vrf myVRF *
```

clear ip rip statistics

Use this command to clear an IPv4 RIP statistics.

Command Syntax

```
clear ip rip statistics (IFNAME |)
```

Parameters

IFNAME	Removes entries from the interface.

Default

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip rip statistics
```

debug rip

Use this command to specify the options for the displayed debugging information for RIP events, RIP packets and RIP NSM.

Use the `no` parameter with this command to disable all debugging.

Command Syntax

```
debug rip (all|events|packet (recv|send|)|packet detail|nsm|rib|bfd|)
no debug rip (all|events|packet (recv|send|)|packet detail|nsm|rib|bfd|)
```

Parameters

`all`	Debug all RIP information.
`bfd`	Debug all RIP and BFD information.
`events`	Debug RIP events.
`nsm`	Debug RIP and NSM communications.
`packet`	Debug RIP packets, only
`recv`	Debug received packets.
`rib`	Debug RIP and RIB communications.
`send`	Debug sent packets.
`detail`	Display detailed information for the sent or received packet.

Default

Disabled

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

The following example specifies the options for the displayed debugging information in `Configure` mode prompt.

```
#configure terminal
(config)#debug rip events
(config)#debug rip packet send detail
(config)#debug rip nsm
```

default-information originate

Use this command to add default routes to the RIP updates.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
default-information originate (always|) (route-map WORD|)
no default-information originate
```

Parameters

`always`	Always advertise default route
`route map`	Route map reference
`WORD`	Pointer to route-map entries

Default

Disabled

Command Mode*

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#default-information originate route-map pmap
```

default-metric

Use this command to specify the metrics to be assigned to redistributed routers.

This command is used in conjunction with the `redistribute` command to make the routing protocol use the specified metric value for all redistributed routes. A default metric is useful in redistributing routes with incompatible metrics. Every protocol has different metrics and can not be compared directly. Default metric provides the standard to compare. All routes that are redistributed will use the default metric.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
default-metric <1-15>
no default-metric
```

Parameter

`<1-15>`	Default metric.

Default

By default, the metric value is set to 1.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example assigns the cost of 10 to the OSPF routes which are redistributed into RIP.

```
#configure terminal
(config)#router rip
(config-router)#redistribute ospf
(config-router)#default-metric 10
```

distance

Use this command to set the administrative distance. The administrative distance is a feature used by the routers to select the path when there are two or more different routes to the same destination from two different routing protocols. A smaller administrative distance indicating a more reliable protocol.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
distance <1-255>
distance <1-255> A.B.C.D/M (WORD|)
no distance
no distance  A.B.C.D/M
```

Parameters

`<1-255>`	Specify the administrative distance value.
`A.B.C.D./M`	Specify the network prefix and length
`WORD`	Specify the access list name.

Default

By default, the administrative distance is 120.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

```
#configure terminal
(config)#router rip
(config-router)#distance 8 10.0.0.0/8 mylist
```

distribute-list

Use this command to filter incoming or outgoing route updates using an access list or a prefix list. You can filter out incoming or outgoing route updates using an access list or a prefix list. If you do not specify the name of the interface, the filter will be applied to all the interfaces.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
distribute-list WORD (in|out) (IFNAME|)
distribute-list prefix WORD (in|out) (IFNAME|)
no distribute-list (in|out) (IFNAME|)
no distribute-list prefix (in|out) (IFNAME|)
```

Parameters

`WORD`	Specify the IPv4 access-list number or name to use.
`prefix`	Filter prefixes in routing updates.
`WORD`	Specify the name of the IPv4 prefix-list to use.
`in`	Filter incoming routing updates.
`out`	Filter outgoing routing updates.
`IFNAME`	Specify the name of the interface on which distribute-list applies.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Example

```
#configure terminal
(config)#router rip
(config-router)#distribute-list prefix myfilter in eth0
```

ip rip authentication key-chain

Use this command to enable RIPv2 authentication on an interface and specify the name of the key chain to be used. If you do not configure a key chain results in no authentication.

Use the `no` parameter with this command to disable this function.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
ip rip authentication key-chain LINE
no ip rip authentication key-chain (LINE|)
```

Parameters

`LINE`	Specify the name of the key chain.

Default

If you do not configure a key chain, authentication is not used.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, interface eth0 is configured key-chain authentication and the name is specified as `mykey`. This name is used to enter the key-chain mode to specify the password. See the key command.

```
#configure terminal
(config)#interface eth0
(config-if)#ip rip authentication key-chain mykey
```

ip rip authentication mode

Use this command to specify the type of authentication mode used for RIP v2 packets.

Use the `no` parameter with this command to restore clear text authentication.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
ip rip authentication mode md5
ip rip authentication mode text
no ip rip authentication mode
```

Parameters

`md5`	Uses the keyed MD5 authentication algorithm.
`text`	Specify the clear text or simple password authentication.

Default

No authentication mode is enabled by default. But, when any authentication key (string or key-chain) is configured, text authentication mode is enabled by default.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows `md5` authentication configured on the `eth1` interface, ensuring authentication of RIP packets received.

```
#configure terminal
(config)#interface eth1
(config-if)#ip rip authentication mode md5
```

ip rip authentication string

Use this command to specify the authentication string or password used by a key.

You can configure authentication for a single key or multiple keys at different times. Use this command to specify password for a single key on an interface.

Use the `no` parameter with this command to disable this feature.

See Appendix 4, *Routing Information Protocol Authentication* for how this command is related to the other authentication commands.

Command Syntax

```
ip rip authentication string LINE
no ip rip authentication string
```

Parameters

`LINE`	Specify the authentication string or password used by a key.

Default

None

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the interface `eth1` is configured to have an authentication string as `guest`, any receiving RIP packet in that interface should have the same string as password.

```
#configure terminal
(config)#interface eth1
(config-if)#ip rip authentication string guest
```

ip rip receive-packet

Use this command to configure the interface to enable the reception of RIP packets.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
ip rip receive-packet
no ip rip receive-packet
```

Parameters

None

Default

Receive-packet is enabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows packet receiving being turned on for interface `eth0`.

```
#configure terminal
(config)#interface eth0
(config-if)#ip rip receive-packet
```

ip rip receive version

Use this command to receive specified version of RIP packets on an interface basis using version control, and override the setting of the version command.

Use the `no` form of this command to use the setting established by the version command.

Command Syntax

```
ip rip receive version (1|2)
ip rip receive version 1 2
no ip rip receive version
```

Parameters

`1`	Specify acceptance of RIP version 1 packets on the interface.
`2`	Specify acceptance of RIP version 2 packets on the interface.
`1 2`	Specify acceptance of RIP version 1 and version 2 packets on the interface.

Default

Version 2

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

In the following example, interface eth1 in configured to receive both RIP version 1 and 2 packets.

```
#configure terminal
(config)#interface eth1
(config-if)#ip rip receive version 1 2
```

ip rip send-packet

Use this command to enable sending RIP packets through the current interface.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
ip rip send-packet
no ip rip send-packet
```

Parameters

None

Default

Send packet is enabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows packet sending being turned on for interface `eth0`.

```
#configure terminal
(config)#interface eth0
(config-if)#ip rip send-packet
```

ip rip send version

Use this command to send RIP packets on an interface using version control. In addition to version 1 and version 2, compatible version packets can be specified. With the parameter 1-compatible, a version 2 RIP interface will broadcast the packets instead of multicasting them.

This command applies to a specific interface and overrides any the version specified by the `version` command.

Use the `no` parameter with this command to use the global RIP version control rules.

Command Syntax

```
ip rip send version (1|2|1-compatible)
ip rip send version 1 2
no ip rip send version
```

Parameters

`1`	Specify sending RIP version 1 packets out of an interface.
`2`	Specify sending RIP version 2 packets out of an interface.
`1 2`	Specify acceptance of RIP version 1 and version 2 packets on the interface.
`1-compatible`	Specify sending RIP version 1 compatible packets from a version 2 RIP interface.

Default

Version 2

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

In the following example, interface eth1 is configured to send both RIP version 1 and 2 packets.

```
#configure terminal
(config)#interface eth1
(config-if)#ip rip send version 1 2
```

ip rip split-horizon

Use this command to perform the split-horizon action on the interface

This command helps avoid including routes in updates sent to the same gateway from which they were learned. Using the `split horizon` command omits routes learned from one neighbor, in updates sent to that neighbor. Using the `poisoned` parameter with this command includes such routes in updates, but sets their metrics to infinity. Thus, advertising that these routes are not reachable.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
ip rip split-horizon
ip rip split-horizon poisoned
no ip rip split-horizon
```

Parameter

`poisoned`	Performs split-horizon with poisoned reverse.

Default

Split horizon poisoned

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ip rip split-horizon poisoned
```

key

Use this command to manage, add or delete authentication keys in a key-chain. This command allows you to enter the Keychain-key mode to set a password for the key.

Use the no option with this command to disable this feature.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
key <0-2147483647>
no key <0-2147483647>
```

Parameters

`<0-2147483647>` Specify a key identifier.

Default

By default, RIP uses level-1-2 if there is no Level-2 instance nor a Level-1-2 instance. Otherwise, it uses level-1.

Command Mode

Keychain mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#key-string prime

(config-keychain)#key 1
(config-keychain-key)#no key-string
```

key chain

Use this command to enter the key chain management mode and to configure a key chain with a key chain name. This command allows you to enter the keychain mode to specify keys on this key chain.

Use the no option with this command to disable this feature.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
key chain WORD
no key chain WORD
```

Parameter

`WORD`	Specify the name of the key chain to manage.

Default

None

Command Mode

Configure mode and Keychain mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the creation of a key chain named `mychain` and the change into `keychain` mode prompt.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#
```

The following example shows the creation of a key chain named `mykeychain3` in the Keychain mode and the addition of an authentication key `key10` in the same mode.

```
(config-keychain)#key chain mykeychain3
(config-keychain)#key 10
(config-keychain-key)#
```

key-string

Use this command to define a password to be used by a key.

Use the `no` parameter with this command to disable this feature.

See Appendix 4, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
key-string LINE
no key-string
```

Parameters

`LINE`	Specify a string of characters to be used as a password by the key.

Default

Disabled

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#key-string prime

(config-keychain)#key 1
(config-keychain-key)#no key-string
```

maximum-prefix

Use this command to configure the maximum prefix.

Use the `no` parameter with this command to disable the limiting of the number of RIP routes in the routing table.

Command Syntax

```
maximum-prefix <1-65535> (<1-100>|)
no maximum-prefix
```

Parameters

`<1-65535>`	The maximum number of RIP routes allowed.
`<1-100>`	Percentage of maximum routes to generate a warning. The default threshold is 75%.

Default

The default maximum-prefix threshold is 75%.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#maximum-prefix 150
```

neighbor

Use this command to specify a neighbor router. It is used for each connected point-to-point link. This command to exchanges non-broadcast routing information. It can be used multiple times for additional neighbors.

`Passive-interface` command disables sending routing updates on an interface. Use the `neighbor` command in conjunction with the `passive-interface` command to send routing updates to specific neighbors.

Use the `no` parameter with this command to disable the specific router.

Command Syntax

```
neighbor A.B.C.D
no neighbor A.B.C.D
```

Parameter

`A.B.C.D`	An IP address of a neighboring router with which the routing information will be exchanged.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#neighbor 10.7.1.12
```

network

Use this command to specify a network as one that runs RIP. This command specifies the networks to which routing updates will be sent and received. If a network is not specified, the interfaces in that network will not be advertised in any RIP update.

Use the `no` parameter with this command to remove the specified network as one that runs RIP.

Command Syntax

```
network A.B.C.D/M
network IFNAME
no network A.B.C.D/M
no network IFNAME
```

Parameters

`A.B.C.D/M`	The IP address prefix and length of this IP network.
`IFNAME`	Alphanumeric string that defines the interface name.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#network 10.0.0.0/8
(config-router)#network eth0
```

offset-list

Use this command to add an offset to in and out metrics to routes learned through RIP. This command specifies the offset value that is added to the routing metric. When the networks match the access list the offset is applied to the metrics. No change occurs if the offset value is zero.

Use the `no` parameter with this command to remove the offset list.

Command Syntax

```
offset-list WORD (in|out) (IFNAME|)
no offset-list WORD (in|out) (IFNAME|)
```

Parameters

`WORD`	Specify the access-list number or names to apply.
`in`	Indicates the access list will be used for metrics of incoming advertised routes.
`out`	Indicates the access list will be used for metrics of outgoing advertised routes.
`IFNAME`	An alphanumeric string that specifies the interface to match.

Default

The default `offset value` is the interface metric value which is defined by the operating system.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

In this example the router examines the RIP updates being sent out from interface eth0 and adds 16 hops to the routes matching the ip addresses specified in the access list `accesslist1`.

```
#configure terminal
(config)#router rip
(config-router)#offset-list accesslist1 in eth0
```

passive-interface

Use this command to block RIP broadcast on the interface.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
passive-interface IFNAME
no passive-interface IFNAME
```

Parameters

`IFNAME`	Specify the interface name.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#passive-interface eth0
```

recv-buffer-size

Use this command to run-time configure the RIP UDP receive-buffer size.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
recv-buffer-size <8192-2147483647>
no recv-buffer-size
```

Parameters

`<8192-2147483647>`

Specify the RIP UDP receive buffer size value.

Default

The default value of the RIP UDP receive-buffer size is 32768.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#recv-buffer-size 150000
```

redistribute

Use this command to redistribute information from other routing protocols.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
redistribute (connected|static|ospf|isis|bgp)
redistribute (connected|static|ospf|isis|bgp) metric <0-16>
redistribute (connected|static|ospf|isis|bgp) route-map WORD
redistribute (connected|static|ospf|isis|bgp) metric <0-16> route-map WORD
no redistribute (connected|static|ospf|isis|bgp)
no redistribute (connected|static|ospf|isis|bgp) metric <0-16>
no redistribute (connected|static|ospf|isis|bgp) route-map WORD
no redistribute (connected|static|ospf|isis|bgp) metric <0-16> route-map WORD
```

Parameters

`bgp`	Redistribute from BGP routes
`connected`	Redistribute from connected routes
`isis`	Redistribute from ISO IS-IS routes
`ospf`	Redistribute from OSPFv3 routes
`static`	Redistribute from static routes
`metric`	Metric value
`<0-16>`	Specify a metric value
`route-map`	Route map reference
`WORD`	Specify name of the route-map

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#redistribute connected


#configure terminal
(config)#router rip
```

```
(config-router)#redistribute connected route-map myroutemap
```

route

Use this command to configure static RIP routes.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
route A.B.C.D/M
no route A.B.C.D/M
```

Parameter

`A.B.C.D/M`	Specify the IP address prefix and length.

Default

No route is added.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

Use this command to add a static RIP route. This command is used most often for debugging purposes and does not show up in the kernel routing table. After adding the RIP route, it can be checked in the RIP routing table.

```
#configure terminal
(config)#router rip
(config-router)#version 1
(config-router)#network 10.10.10.0/24
(config-router)#network 10.10.11.0/24
(config-router)#neighbor 10.10.10.10
(config-router)#route 10.10.10.0/24

(config-router)#version 1
(config-router)#network 10.10.10.0/24
(config-router)#network 10.10.11.0/24
(config-router)#no route 10.10.10.0/24
```

router rip

Use this global command to enable a RIP routing process.

Use the `no` parameter with this command to disable RIP routing.

Command Syntax

```
router rip
no router rip
```

Parameter

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This command is used to begin the RIP routing process.

```
#configure terminal
(config)#router rip
(config-router)#version 1
(config-router)#network 10.10.10.0/24
(config-router)#network 10.10.11.0/24
(config-router)#neighbor 10.10.10.10
```

send-lifetime

Use this command to specify the time period during which the authentication key on a key chain can be sent.

Use the `no` parameter with this command to negate this command.

Command Syntax

```
send-lifetime HH:MM:SS <1-31> MONTH <1993-2035> HH:MM:SS <1-31> MONTH <1993-2035>
send-lifetime HH:MM:SS <1-31> MONTH <1993-2035> HH:MM:SS MONTH <1-31> <1993-2035>
send-lifetime HH:MM:SS MONTH <1-31> <1993-2035> HH:MM:SS <1-31> MONTH <1993-2035>
send-lifetime HH:MM:SS MONTH <1-31> <1993-2035> HH:MM:SS MONTH <1-31> <1993-2035>
send-lifetime HH:MM:SS <1-31> MONTH <1993-2035> infinite
send-lifetime HH:MM:SS MONTH <1-31> <1993-2035> infinite
send-lifetime HH:MM:SS <1-31> MONTH <1993-2035> duration <1-2147483646>
send-lifetime HH:MM:SS MONTH <1-31> <1993-2035> duration <1-2147483646>
no send-lifetime
```

Parameters

`HH:MM:SS`	Specify the start time of send-lifetime in hours, minutes and seconds.
`<1-31>`	Specify the day of the month to start.
`MONTH`	Specify the month of the year to start (the first three letters of the month, for example, `Jan.`).
`<1993-2035>`	Specify the year to start.
`HH:MM:SS`	Specify the time when send-lifetime expires in hours, minutes and seconds.
`<1-31>`	Specify the day of the month to expire.
`MONTH`	Specify the month of the year to expire (the first three letters of the month, for example, `Jan.`).
`<1993-2035>`	Specify the year to expire.
`duration`	Specify the duration of the key in seconds <1-2147483646>.
`infinite`	Specify the end time to never expire.

Default

Disabled

Command Mode

Keychain-key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of `send-lifetime` for `key 1` on the key chain named `mychain`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key 1
(config-keychain-key)#send-lifetime 03:03:01 Jan 3 2004 04:04:02 Dec 6 2006
```

show debugging rip

Use this command to display the RIP debugging status for these debugging options: nsm debugging, RIP event debugging, RIP packet debugging and RIP nsm debugging.

Command Syntax

```
show debugging rip
```

Parameters

None

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging rip
R2#show debugging rip
RIP debugging status:
  RIP event debugging is on
  RIP packet detail debugging is on
  RIP RIB debugging is on
  RIP NSM debugging is on
  RIP BFD debugging is on

R2#
```

show ip protocols rip

Use this command to display RIP process parameters and statistics.

Command Syntax

```
show ip protocols
show ip protocols rip
```

Parameters

None

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is an example of the output from the `show ip protocols rip` command:

```
#show ip protocols rip
Routing Protocol is "rip"
Sending updates every 30 seconds with +/-50%, next due in 12 seconds
Timeout after 180 seconds, garbage collect after 120 seconds
Outgoing update filter list for all interface is not set
Incoming update filter list for all interface is not set
Default redistribution metric is 1
Redistributing: connected static
Default version control: send version 2, receive version 2
Interface        Send  Recv   Key-chain
     eth0                2   2
Routing for Networks:
  10.10.0.0/24
Routing Information Sources:
  Gateway          BadPackets BadRoutes  Distance Last Update
Distance: (default is 120)
#
```

Figure 1-41 Explains the show command output details.

Table 1-41: Show ip protocols output details

Field	Description
Routing Protocol is "rip"	Specifies the routing protocol used.
Sending updates every 30 seconds	Specifies the time between sending updates.

Next due in 12 seconds	Precisely when the next update is due to be sent.
Timeout after 180 seconds	Specifies the value of the timeout parameter.
Redistributing	Lists the protocol that is being redistributed.
Routing for Networks	Specifies the networks for which the routing process is currently injecting routes.
Routing Information Sources	Lists all the routing sources the IP Infusion software is using to build its routing table.

show ip rip

Use this command to show RIP routes.

Command Syntax

```
show ip rip (database|)
```

Parameters

database	Specify to display information about the IP RIP database.

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following output displays the RIP routing table with the destination network, nexthop and metric to reach it.

```
#show ip rip
Codes: R - RIP, K - Kernel, C - Connected, S - Static, O - OSPF, I - IS-IS,
B - BGP
Network Next Hop Metric From If Time
K 0.0.0.0/0 10.0.1.1 16 eth1 01:58
C 10.0.1.0/24 1 eth1
S 10.10.10.0/24 1 eth0
C 10.10.11.0/24 1 eth0
S 192.168.101.0/24 1 eth0
R 192.192.192.0/24 1 --
```

Table 1-42 shows the status codes displayed at the start of a route entry.

Table 1-42: Status codes

Status Code	Meaning	Description
R	RIP	RIP prevents routing loops by implementing a limit on the number of hops allowed in a path from source to destination.
K	Kernel	Kernel is central component of operating system.
C	Connected	Redistribute from locally connected networks.
S	Static	Connections in a static network are fixed links, while connections in a dynamic network are established on the fly as needed.
O	OSPF	Open Shortest Path First (OSPF) is a routing protocol for Internet Protocol (IP) networks.

I	IS-IS	Intermediate System to Intermediate System (IS-IS) is a routing protocol designed to move information efficiently within a host network.
B	BGP	BGP makes routing decisions based on paths, rules or network policies configured by a network administrator.

show ip rip interface

Use this command to display information about RIP interfaces. You can specify an interface name to display information about a specific interface.

Command Syntax

```
show ip rip interface (IFNAME|)
```

Parameters

IFNAME	Name of the interface for which information is to be displayed.

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following output displays the RIP routing table with the destination network, nexthop and metric to reach it.

```
#show ip rip interface
lo is up, line protocol is up
RIP is not enabled on this interface
eth0 is up, line protocol is up
RIP is not enabled on this interface
eth1 is down, line protocol is down
RIP is not enabled on this interface
eth2 is up, line protocol is up
Routing Protocol: RIP
Receive RIP packets
Send RIPv1 Compatible
Passive interface: Disabled
Split horizon: Enabled with Poisoned Reversed
IP interface address:
10.10.1.1/24
10.10.2.1/24
```

Figure 1-43 Explains the show command output details.

Table 1-43: Show ip rip interface output details

Field	Description
Network	IP address of a network entity.
Nexthop	IP address of the next system that is used when forwarding a packet to the destination network.

Metric	If shown, the value of the inter-autonomous system metric.
Routing Protocol	Specifies the routing protocol used.
Passive Interface	Used in all routing protocols to disable sending updates out from a specific interface.
Split horizon	the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.
IP Interface address	IP address of the RIP peer neighbor.

show ip rip statistics

Use this command to display information about RIP statistics. You can specify an interface name to display information about a specific interface.

Command Syntax

```
show ip rip statistics (IFNAME|)
```

Parameters

IFNAME	Name of the interface for which information is to be displayed.

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following output displays the RIP routing table with the destination network, nexthop and metric to reach it.

```
#show ip rip statistics eth1
Interface Name : eth1
  Sent Multicast Updates  : 3
  Sent Multicast Requests : 1
  Sent Unicast Updates    : 0
  Sent Unicast Requests   : 0
  Recv Multicast Updates  : 3
  Recv Multicast Requests : 0
  Recv Unicast Updates    : 1
  Recv Unicast Requests   : 0
  Recv Bad Packets        : 0
  Recv Bad Routes         : 0
```

Figure 1-44 Explains the show command output details.

Table 1-44: Show ip rip statistics output details

Field	Description
Network	IP address of a network entity.
Nexthop	IP address of the next system that is used when forwarding a packet to the destination network.
Metric	If shown, the value of the inter autonomous system metric.
Sent updates	Number of RIP routing updates that have been sent on Multicast/Unicast interface.

Sent Request	Number of RIP routing request that have been sent on Multicast/Unicast interface.
Recv updates	Number of RIP routing updates that have been received on Multicast/Unicast interface.
Recv Request	Number of RIP routing request that have been received on Multicast/Unicast interface.
Recv Bad Packets	Number of packets that were received on this interface and were not processed for any reason.
Recv Bad Routes	Number of route entries that were received on this interface and were not processed for any reason.

snmp restart rip

Use this command to restart SNMP in Routing Information Protocol (RIP)

Command Syntax

```
snmp restart rip
```

Parameters

None

Default

By default, snmp restart is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#snmp restart rip
```

timers basic

Use this command to adjust routing network timers.

This command adjusts the RIP timing parameters. Every 30 seconds, an update is sent out containing the complete routing table to every neighboring router. When the time specified by the timeout parameter expires, the route is no longer valid. However, it is retained in the routing table for a short time so that neighbors are notified that the route has been dropped. When the time specified by the garbage parameter expires, the route is finally removed from the routing table. Until the garbage time expires, the route is included in all updates sent by the router.

All routers in the network must have the same timers to allow RIP to execute a distributed and asynchronous routing algorithms. The timers should not be synchronized as it might lead to unnecessary collisions on the network.

Use the `no` parameter with this command to restore the default routing network timers.

Command Syntax

```
timers basic <5-2147483647> <5-2147483647> <5-2147483647>
no timers basic
```

Parameters

`<5-2147483647>`	Specify the routing table update timer in seconds. The default is 30 seconds.
`<5-2147483647>`	Specify the routing information timeout timer in seconds. The default is 180 seconds. After this interval has elapsed and no updates for a route are received, the route is declared invalid.
`<5-2147483647>`	Specify the routing garbage collection timer in seconds. The default is 120 seconds.

Default

The default routing table update time is 30 seconds.

The default routing information timeout time is 180 seconds.

The default routing garbage collection time is 120 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#timers basic 30 180 120

(config)#router rip
(config-router)#no timers basic
```

version

Use this command to specify a RIP version used globally by the router. RIP can be run in version 1 as well as version 2 mode. Version 2 has more features than version 1 including authentication. Once the rip version is set, rip packets of that version will be received and sent on all the rip-enabled interfaces.

Use the `no` parameter with this command to restore the default version.

Note: The `ip rip receive version` command and the `ip rip send version` command override the value set by the `version` command.

Command Syntax

```
version <1-2>
no version
```

Parameters

`<1-2>`	Specify the version of RIP processing.

Default

Version 2

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rip
(config-router)#version 1
(config-router)#network 10.10.10.0/24
(config-router)# network 10.10.11.0/24
```

CHAPTER 2 RIPng Commands

This chapter provides an alphabetized reference for each of the Routing Information Protocol next generation (RIPng) commands, which support IPv6. It includes the following commands:

- aggregate-address
- cisco-metric-behavior
- clear ipv6 rip route
- debug ipv6 rip
- default-information originate
- default-metric
- distance
- distribute-list
- ipv6 rip metric-offset
- ipv6 rip split-horizon
- ipv6 router rip
- neighbor
- offset-list
- passive-interface
- recv-buffer-size
- redistribute
- route
- route-map
- router ipv6 rip
- show debugging ipv6 rip
- show ipv6 protocols rip
- show ipv6 rip
- show ipv6 rip interface
- timers basic

© 2025 IP Infusion Inc. Proprietary

aggregate-address

Use this command to set an aggregate RIPng route announcement.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
aggregate-address X:X::X:X/M
no aggregate-address X:X::X:X/M
```

Parameter

`X:X::X:X/M`	Specify an aggregate network (IPv6 address prefix and length).

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#aggregate-address 3ffe:8088::/32

(config)#router ipv6 rip
(config-router)#no aggregate-address 3ffe:8088::/32
```

cisco-metric-behavior

Use this command to enable or disable the metric update as Cisco.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
cisco-metric-behavior (enable|disable)
no cisco-metric-behavior
```

Parameters

`enable`	Enable updating the metric consistent with Cisco.
`disable`	Disable updating the metric consistent with Cisco.

Default

By default, the Cisco metric-behavior is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to enable the metric update behavior to be consistent with Cisco in the Router mode.

```
#configure terminal
(config)#router ipv6 rip
(config-router)#cisco-metric-behavior enable
```

clear ipv6 rip route

Use this command to clear specific data from the RIPng routing table.

Command Syntax

```
clear ipv6 rip route (X:X::X:X/M|rip|kernel|connected|static|ospf6|isis|bgp|all)
```

Parameters

`X:X::X:X/M`	Removes entries which exactly match this destination address from the RIPng routing table.
`bgp`	Removes only BGP routes from the RIP routing table.
`connected`	Removes entries for connected routes from the RIP routing table.
`isis`	Removes only IS-IS routes from the RIP routing table
`kernel`	Removes kernel entries from the RIP routing table.
`ospf`	Removes only OSPF routes from the RIP routing table.
`static`	Removes static entries from the RIP routing table.
`all`	Removes the entire RIP routing table.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ipv6 rip route isis
#clear ipv6 rip route 3ffe:ffff::/16
```

debug ipv6 rip

Use this command to specify the options for the displayed debugging information for RIPng events, RIPng packets and RIPng NSM communications.

Use the `no` option with this command to turn off debugging options for RIPng.

Command Syntax

```
debug ipv6 rip (all|events|packet send|packet recv|packet|packet detail|nsm|rib|)
no debug ipv6 rip (all|events|packet send|packet recv|packet|packet
  detail|nsm|rib|)
```

Parameters

`all`	Debug all RIP information.
`events`	Debug RIP events.
`nsm`	Debug RIP and NSM communications.
`rib`	Debug RIP and RIB communications.
`packet`	Debug RIP packets, only Routing Information Protocol
`recv`	Debug received packets.
`send`	Debug sent packets.
`detail`	Display detailed information for the sent or received packet.

Default

Disabled

Command Mode

Privileged Exec mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example specifies the options for the displayed debugging information in `Configure` mode prompt.

```
#configure terminal
(config)#debug ipv6 rip events
(config)#debug ipv6 rip packet send detail
(config)#debug ipv6 rip nsm
```

default-information originate

Use this command to generate a default route into the RIPng.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
default-information originate
no default-information originate
```

Parameters

None

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#default-information originate
```

default-metric

Use this command to specify the metrics to be assigned to redistributed routes.

Use the `no` parameter with this command to disable this feature.

For more details about this command, see the IPv4 version of this command (default-metric).

Command Syntax

```
default-metric <1-15>
no default-metric
```

Parameter

`<1-15>`	Specify the default metric.

Default

By default, the metric value is set to 1.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#default-metric 8
```

distance

Use this command to set the administrative distance for RIP.

Use the no option with this command to disable this function.

For more details about this command, see the IPv4 version of this command (distance).

Command Syntax

```
distance <1-255>
no distance
```

Parameter

<1-255>	Specify the administrative distance value.

Default

By default, the administrative distance is 120.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 rip
(config-router)#distance 100
```

distribute-list

Use this command to filter incoming or outgoing route updates using the access-list or the prefix-list. You can filter out incoming or outgoing route updates using access-list or prefix-list. If you do not specify the name of the interface, the filter will be applied to all the interfaces.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
distribute-list WORD (in|out) (IFNAME|)
distribute-list prefix WORD (in|out) (IFNAME|)
no distribute-list (in|out) (IFNAME|)
no distribute-list prefix (in|out) (IFNAME|)
```

Parameters

`WORD`	Specify the IPv6 access-list number or name to use.
`prefix`	Filter prefixes in routing updates.
`WORD`	Specify the name of the IPv6 prefix-list to use.
`in`	Filter incoming routing updates.
`out`	Filter outgoing routing updates.
`IFNAME`	Specify the name of the interface on which distribute-list applies.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Example

```
#configure terminal
(config)#router ipv6 rip
(config-router)#distribute-list prefix myfilter in eth0
```

ipv6 rip metric-offset

Use this command to set RIP metric offset.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
ipv6 rip metric-offset <1-16>
no ipv6 rip metric-offset
```

Parameter

`<1-16>`	Set a metric value

Default

None

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 rip metric-offset 1

(config)#interface eth0
(config-if)#no ipv6 rip metric-offset
```

ipv6 rip split-horizon

Use this command to perform the split-horizon action on the interface.

Use the `no` parameter with this command to disable this function.

For more details about this command, see the IPv4 version of this command (ip rip split-horizon).

Command Syntax

```
ipv6 rip split-horizon
ipv6 rip split-horizon poisoned
no ipv6 rip split-horizon
```

Parameter

`poisoned`	Performs split-horizon with poisoned reverse.

Default

By default, Split horizon poisoned is enabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth1
(config-if)#ipv6 rip split-horizon

(config)#interface eth1
(config-if)#no ipv6 rip split-horizon
```

ipv6 router rip

Use this command to enable RIPng routing on the interface.

Use the `no` parameter with this command to disable RIPng routing.

Command Syntax

```
ipv6 router rip
no ipv6 router rip
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 router rip
```

neighbor

Use this command to specify a neighbor router.

Use the `no` parameter with this command to disable the specific router.

For more details about this command, see the IPv4 version of this command (neighbor).

Command Syntax

```
neighbor X:X::X:X IFNAME
no neighbor X:X::X:X IFNAME
```

Parameters

`X:X::X:X`	Specify a link-local IP address of a neighboring router with which the routing information is exchanged.
`IFNAME`	Specify the name of the interface.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ipv6 rip
(config-router)#neighbor 80::1 eth0
```

offset-list

Use this command to add an offset to in and out metrics to routes learned through RIPng.

Use the `no` parameter with this command to remove this function.

For more details about this command, see the IPv4 version of this command (offset-list).

Command Syntax

```
offset-list WORD (in|out) (IFNAME|)
no offset-list in|out (IFNAME|)
```

Parameters

`WORD`	Specify the access-list number or names to apply.
`in`	Indicates the access list will be used for metrics of incoming advertised routes.
`out`	Indicates the access list will be used for metrics of outgoing advertised routes.
`IFNAME`	An alphanumeric string that specifies the interface to match.

Default

The default offset value is the metric value of the interface which is defined by the operating system.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.1.0.

Examples

In this example the router examines the RIP updates being sent out from interface eth0 and adds 16 hops to the routes matching the ip addresses specified in the access list `accesslist1`.

```
#configure terminal
(config)#router ipv6 rip
(config-router)#offset-list accesslist1 in eth0
```

passive-interface

Use this command to suppress routing updates on an interface.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
passive-interface IFNAME
no passive-interface IFNAME
```

Parameters

`IFNAME`	Specify the interface name.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#passive-interface eth0
```

recv-buffer-size

Use this command to run-time configure the RIPng UDP receive-buffer size.

Use the `no` parameter with this command to return to the default value.

Command Syntax

```
recv-buffer-size <8192-2147483647>
no recv-buffer-size
```

Parameters

`<8192-2147483647>`

Specify the RIP UDP receive buffer size value.

Default

The default value of the RIP UDP receive-buffer size is 8192.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#recv-buffer-size 150000
```

redistribute

Use this command to redistribute information from other routing protocols.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
redistribute (connected|static|ospf|isis|bgp)
redistribute (connected|static|ospf|isis|bgp) metric <0-16>
redistribute (connected|static|ospf|isis|bgp) route-map WORD
redistribute (connected|static|ospf|isis|bgp) metric <0-16> route-map WORD
no redistribute (connected|static|ospf|isis|bgp)
no redistribute (connected|static|ospf|isis|bgp) metric <0-16>
no redistribute (connected|static|ospf|isis|bgp) route-map WORD
no redistribute (connected|static|ospf|isis|bgp) metric <0-16> route-map WORD
```

Parameters

`bgp`	Redistribute from BGP routes
`connected`	Redistribute from connected routes
`isis`	Redistribute from ISO IS-IS routes
`ospf`	Redistribute from OSPF routes (version 3)
`static`	Redistribute from static routes
`metric`	Metric value
`<0-16>`	Specify a metric value
`route-map`	Route map reference
`WORD`	Specify name of the route-map

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#redistribute connected route-map mymap
(config-router)#redistribute connected metric 8
```

route

Use this command to debug the specified route advertisement. Use this command to configure static RIPng routes.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
route X:X::X:X/M
no route X:X::X:X/M
```

Parameter

`X:X::X:X/M`	Specify the IPv6 address prefix and length.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#route 3ffe:1234:5678::1/64
```

route-map

Use this command to set a route map for input or output filtering on a specified interface.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
route-map WORD (in|out) IFNAME
no route-map (in|out) IFNAME
```

Parameters

`WORD`	Specify a route map name
`in`	Specify to set the route map for input filtering
`out`	Specify to set the route map for output filtering
`IFNAME`	Specify an interface name to which to associate the route map

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#route-map myRM in eth1
```

router ipv6 rip

Use this global command to enable a RIPng routing process.

Use the `no` parameter with this command to disable the RIPng routing process.

Command Syntax

```
router ipv6 rip
no router ipv6 rip
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#
```

show debugging ipv6 rip

Use this command to display the RIPng debugging status for RIPng NSM, RIPng events, and RIPng packets.

Command Syntax

```
show debugging ipv6 rip
```

Parameters

None

Command Mode

Exec Mode and Privileged Exec Mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging ipv6 rip
 RIPng packet debugging is on
```

show ipv6 protocols rip

Use this command to display RIPng process parameters and statistics.

Command Syntax

```
show ipv6 protocols rip
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ipv6 protocols rip` command.

```
#show ipv6 protocols rip
Routing Protocol is "ripng"
  Sending updates every 30 seconds with +/-50%, next due in 10 seconds
  Timeout after 180 seconds, garbage collect after 120 seconds
  Outgoing update filter list for all interface is not set
  Incoming update filter list for all interface is not set
  Default redistribute metric is 1
  Redistributing: connected
  Routing for Networks:
    3ffe:1::/64
#
```

show ipv6 rip

Use this command to show RIP routes.

Command Syntax

```
show ipv6 rip (database|)
```

Parameters

`database`	Specify to display information about the IPv6 RIP database.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the show ipv6 rip database command.

```
#show ipv6 rip database
Codes: R - RIP, K - Kernel, C - Connected, S - Static, O - OSPF, I - IS-IS,
B - BGP, a - aggregate, s - suppressed
Network Next Hop If Met Tag Time
R 3ffe:1234:5678::/64 fe80::3 eth1 3 0 02:28
C 3ffe:ffff:1::/64 :: eth0 1 0
Ra 3ffe:ffff:2::/48 -- 1 0
Rs 3ffe:ffff:2::/48 fe80::3 eth1 3 0 02:32
Cs 3ffe:ffff:2::/64 :: eth1 1 0
R 3ffe:ffff:ffff:ffff::/64 fe80::3 eth1 3 0 02:28
```

show ipv6 rip interface

Use this command to display information about the RIPng interfaces. You can specify an interface name to display information about a specific interface.

Command Syntax

```
show ipv6 rip interface (IFNAME|)
```

Parameters

IFNAME	Name of the interface for which information is to be displayed.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ipv6 rip interface` command.

```
#show ipv6 rip interface
lo is up, line protocol is up
RIPng is not enabled on this interface
eth0 is up, line protocol is up
RIPng is not enabled on this interface
eth1 is down, line protocol is down
RIPng is not enabled on this interface
eth2 is up, line protocol is up
Routing Protocol: RIPng
Passive interface: Disabled
Split horizon: Enabled with Poisoned Reversed
IP interface address:
3ffe:ffff::1/64
3ffe:fffe::1/64
```

timers basic

Use this command to adjust routing network timers.

Use the `no` parameter with this command to restore the defaults.

For more details about this command, see the IPv4 version of this command (timers basic).

Command Syntax

```
timers basic <5-2147483647> <5-2147483647> <5-2147483647>
no timers basic
```

Parameters

`<5-2147483647>`	Specify the routing table update timer in seconds. The default is 30 seconds.
`<5-2147483647>`	Specify the routing information timeout timer in seconds. The default is 180 seconds. After this interval has elapsed and no updates for a route are received, the route is declared invalid.
`<5-2147483647>`	Specify the routing garbage collection timer in seconds. The default is 120 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 rip
(config-router)#timers basic 30 180 120

(config)#router ipv6 rip
(config-router)#no timers basic
```

CHAPTER 3 Routing Information Protocol VPN Commands

This chapter provides information about RIP VPN commands. These commands are available when the RIP Provider Edge (PE) and Customer Edge (CE) feature is supported. Using these commands, VPN you can use RIP to receive information which the CE-router places into the connected Virtual Routing and Forwarding (VRF) from the receiving interface. The information is then advertised across the VPN backbone between PE-routers.

To provide a VPN service, the PE-router needs to be configured so that any routing information learned from a VPN customer interface can be associated with a particular VRF. This is achieved using any standard routing protocol process (RIP, OSPF, BGP or static routes etc).

This chapter contains the following commands:

- show ip rip interface vrf
- show ip rip vrf
- show ip vrf

show ip rip interface vrf

Use this command to display VRF information. This command is supported in RIP (IPv4).

Command Syntax

```
show ip rip interface vrf WORD (IFNAME|)
```

Parameters

WORD	Specify the name for the VRF instance.
IFNAME	Specify name for the interface.

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip rip interface vrf myVRF

eth1 is up, line protocol is up
  Routing Protocol: RIP
    VPN Routing/Forwarding: myVRF
    Receive RIP packets
    Send RIP packets
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IP interface address:
      1.1.1.92/24
eth3 is up, line protocol is up
  RIP is not enabled on this interface
```

Figure 3-45 Explains the show command output details.

Table 3-45: Show ip rip interface vrf output details

Field	Description
Routing Protocol	Specifies the routing protocol used.
VPN Routing/Forwarding	Specifies the name of the virtual routing and forwarding (VRF) instance.
Receive RIP packets	Number of RIP packets that were received on this interface.
Send RIP packets	Number of RIP packets that were send on this interface.
Passive Interface	Used in all routing protocols to disable sending updates out from a specific interface.

Split horizon	the routing switch does not advertise a route on the same interface as the one on which the routing switch learned the route.
IP Interface address	IP address of the RIP peer neighbor.

show ip rip vrf

Use this command to display VRF information. This command is supported in RIP (IPv4).

Command Syntax

```
show ip rip (database) vrf WORD
```

Parameters

`database`	Specify to display information about the IP RIP database.
`WORD`	Specify the name for the VRF instance.

Default

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip rip database vrf myVRF
Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP

   Network            Next Hop         Metric From            If     Time
Rc 1.1.1.0/24                                1                eth1
S  72.72.75.0/24      98.98.8.2              1                eth3
```

Table 3-46 shows the status codes displayed at the start of a route entry.

Table 3-46: Status codes

Status Code	Meaning	Description
R	RIP	RIP prevents routing loops by implementing a limit on the number of hops allowed in a path from source to destination.
Rc	RIP Connected	Redistribute from locally RIP connected networks.
Rs	RIP Static	Connections in a RIP static network are fixed links, while connections in a dynamic network are established on the fly as needed.
K	Kernel	Kernel is central component of operating system.
C	Connected	Redistribute from locally connected networks.

Rs	RIP Static	Connections in a RIP static network are fixed links, while connections in a dynamic network are established on the fly as needed.
O	OSPF	Open Shortest Path First (OSPF) is a routing protocol for Internet Protocol (IP) networks.
I	IS-IS	Intermediate System to Intermediate System (IS-IS) is a routing protocol designed to move information efficiently within a host network.
B	BGP	BGP makes routing decisions based on paths, rules or network policies configured by a network administrator.
	Network	IP address of a network entity.
	Nexthop	IP address of the next system that is used when forwarding a packet to the destination network.
	Metric	If shown, the value of the inter-autonomous system metric.

Figure 3-47 Explains the show command output details.

Table 3-47: Show ip rip vrf output details

Field	Description
Network	IP address of a network entity.
Nexthop	IP address of the next system that is used when forwarding a packet to the destination network.
Metric From	If shown, the value of the inter-autonomous system metric.
If	Ethernet interface.
Time	Specifies the time between forwarding packets.

show ip vrf

Use this command to display VRF information. This command is supported in RIP (IPv4).

Command Syntax

```
show ip vrf
show ip vrf WORD
```

Parameters

WORD	Specify the name for the VRF instance.

Default

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip vrf myVRF

VRF myVRF, FIB ID 1
Router ID: 1.1.1.2 (config)
Interfaces:
  eth1
  eth3
VRF myVRF; (id=1); RIP enabled Interfaces:
  eth1
```

CHAPTER 4 Routing Information Protocol Authentication

To support RIPv2 message authentication, you can choose plain text or MD5 authentication, with the option for a single key or multiple keys in different modes and stages.

Single Key Authentication

Use the following steps to configure route to enable RIPv2 authentication using a single key or password:

1. Define the authentication string or password

   In the Interface mode, specify the authentication string or password used by the key using the following command:

   ```
   ip rip authentication string LINE
   ```

   where `LINE` is the authentication string or password

2. Specify mode of authentication for the interface

   In the Interface mode, specify either text or MD5 authentication using the following command:

   ```
   ip rip authentication mode md5|text
   ```

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ip rip authentication string mykey
(config-if)#ip rip authentication mode md5
```

Multiple Keys Authentication

Use the following steps to configure route to enable RIPv2 authentication using multiple keys at different times:

1. Define a key chain

   In the Configure mode, identify a key chain with a key chain name using the following command:

   ```
   key chain KEYNAME
   ```

   where `KEYNAME` is the name of the chain to manage.

2. Define the key(s)

   In the Keychain mode, specify a key on this key chain using the following command:

   ```
   key KEYID
   ```

   where `KEYID` = `<0-2147483647>` Key Identifier number

3. Define the authentication string or password

   In the Keychain-key mode, define the password used by a key, using the following command:

   ```
   key-string LINE
   ```

   where `LINE` is a string of characters to be used as a password by the key.

4. Set key management options

   This step can be performed at this stage or later when multiple keys are used. The options are configured in the `keychain-key` command mode.

- Set the time period during which the authentication key on a key chain is received as valid, using the following command:

  ```
  accept-lifetime START END
  ```

  where `START` and `END` are the beginning and end of the time period.

- Set the time period during which the authentication key on a key chain can be sent using the following command:

  ```
  send-lifetime START END
  ```

  where `START` and `END` are the beginning and end of the time period.

5. Enable authentication on an interface

   In the Interface mode, enable authentication on an interface and specify the key chain to be used, using the following command:

   ```
   ip rip authentication key-chain CHAINNAME
   ```

   where `CHAINNAME` is a set of valid authentication keys

6. Specify mode of authentication for the interface

   In the Interface mode, specify either text or MD5 authentication using the following command:

   ```
   ip rip authentication mode md5|text
   ```

Example

In the following example, a password `toyota` is set for a `key 1` in a key chain `cars`. On Interface `eth0` authentication is enabled and the authentication mode is set as `MD5`.

```
#configure terminal
(config)#key chain cars
(config-keychain)#key 1
(config-keychain-key)#key-string toyota
(config-keychain-key)#accept-lifetime 10:00:00 Oct 08 2002 duration 43200
(config-keychain-key)#send-lifetime 10:00:00 Oct 8 2002 duration 43200
(config-keychain-key)#exit
(config-keychain)#exit
(config)#interface eth0
(config-if)#ip rip authentication key-chain cars
(config-if)#ip rip authentication mode md5
(config-if)#exit
```

VRF Lite Configuration

CHAPTER 1 VRF Configuration

Overview

Virtual routing and forwarding (VRF) is a technology that allows multiple instances of a routing table to co-exist within the same router at the same time. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other. VRF may be implemented in a network device by distinct routing tables known as forwarding information bases – one per routing instance.

Topology

Figure 1-68: Device topology

Default VRF

`#con t`	Enter the router configuration mode
`(config)#interface eth1`	Switch to interface eth1
`(config-if)#ip address 3.3.3.2/24`	Configure the ip address **`3.3.3.2`** to eth1
`(config-if)#exit`	Exit interface mode.

Adding a Static Route

`#con t`	Enter the router configuration mode
`(config)#ip route  20.20.20.0/24 eth1`	Add static route with eth1 as exit interface

User-Defined VRF

`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1
`(config)#exit`	Exit configure mode.
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf1`	Associate eth1 to vrf1.
`(config-if)#ip address 3.3.3.2/24`	Configure the IP address 3.3.3.2 to eth1
`(config-if)#exit`	Exit interface mode.

© 2025 IP Infusion Inc. Proprietary

Adding a Static Route

`#con t`	Enter the router configuration mode
`(config)# ip route vrf vrf1 20.20.20.0/24 eth1`	Add static route in vrf1 with eth1 as exit interface

Validation

```
#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     127.0.0.0/8 is directly connected, lo, 00:14:59
C     192.168.52.0/24 is directly connected, eth0, 00:14:55
IP Route Table for VRF "management"
IP Route Table for VRF "vrf1"
C     3.3.3.0/24 is directly connected, eth1, 00:00:44
S     20.20.20.0/24 [1/0] is directly connected, eth1, 00:00:08

Gateway of last resort is not set
```

CHAPTER 2 OSPF Configuration

Overview

Open Shortest Path First (OSPF) is an interior routing protocol operating within a single autonomous system (AS) that uses a link state routing algorithm. OSPF gathers link state information from available routers and constructs a topology map of the network. The topology determines the routing table presented to the Internet layer which makes routing decisions based solely on the destination IP address in IP packets.

This chapter covers OSPF configuration in non-default VRF.

Topology

Figure 2-69: OSPF topology for VRF

Configuration IPv4 VRF

R1

`#con t`	Enter the router config mode.
`(config)#ip vrf vrf1`	Create vrf1
`((config-vrf)#exit`	Exit VRF mode
`(config)#router ospf 1 vrf1`	Associate the ospf process with vrf1.
`(config-router)#network 2.2.2.0/24 area 0`	Specify the network type and area 0.
`(config-router)#ex`	Exit the OSPF configuration mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip vrf forwarding vrf1`	Associate eth1 to vrf1.
`(config-if)#ip address 2.2.2.1/24`	Assign the IP address 2.2.2.1 to eth1 in vrf1

R2

`#con t`	Enter the router config mode
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#exit`	Exit VRF mode
`(config)#router ospf 1 vrf1`	Associate the ospf process with vrf1
`(config-router)#network 2.2.2.0/24 area 0`	Specify the network type and area 0.
`(config-router)#ex`	Exit router mode.
`(config)#interface eth1`	Enter interface mode.

(config-if)#ip vrf forwarding vrf1	Associate eth1 to vrf1.
(config-if)#ip address 2.2.2.2/24	Assign the IP address 2.2.2.1 to eth1 in vrf1

Validation

R1

```
#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 1 VRF(vrf1):
Neighbor ID     Pri   State           Dead Time   Address       Interface        Instance
ID
2.2.2.2    1     Full/Backup        00:00:30        2.2.2.2     eth1              0

#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "vrf1"
C     2.2.2.0/24 is directly connected, eth1, 00:11:31

Gateway of last resort is not set
```

R2

```
#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 1 VRF(vrf1):
Neighbor ID     Pri   State         Dead Time   Address       Interface          Instance
ID
2.2.2.1   1       Full/Backup      00:00:35        2.2.2.1        eth1              0

#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "vrf1"
C     2.2.2.0/24 is directly connected, eth1, 00:11:31
Gateway of last resort is not set
```

CHAPTER 3 RIP Configuration

The Routing Information Protocol (RIP) is a distance-vector routing protocol which uses the hop count as a routing metric. RIP prevents routing loops by limiting the number of hops allowed — (15) in a path from the source to a destination. This hop limit, however, also limits the size of networks that RIP can support. A hop count of 16 is considered an infinite distance and used to indicate inaccessible, inoperable, or otherwise undesirable routes in the selection process.

Note: This chapter covers RIP configuration in non-default VR and non-default VRF.

Topology

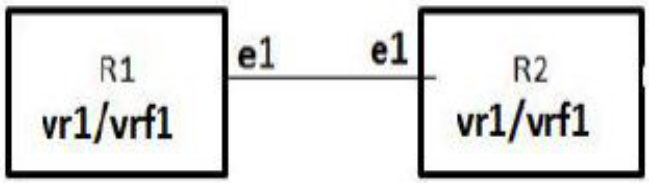

Figure 3-70: RIP topology for VR/VRF

Configuration IPv4 VRF

R1

`#configure terminal`	Enter configure mode.
`(config)#virtual-router VR1`	Create virtual router VR1.
`(config-vr)#load rip`	Load the RIP module in VR1.
`(config)#exit`	Exit VR mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#virtual-router forwarding VR1`	Associate eth1 to VR1.
`(config-if)#exit`	Exit interface mode.
`(config)#exit`	Exit configure mode.
`#login virtual-router VR1`	Log in to virtual-router VR1.
`>en`	Enter privileged exec mode.
`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1
`((config-vrf)#exit`	Exit VRF mode
`(config)#router rip`	Enter router mode.
`(config-router)#version 2`	Specify RIP version 2.
`(config-router)#address-family ipv4 vrf  vrf1`	Enter address family mode for vrf1
`(config-router-af)#network 2.2.2.0/24`	Advertise the connected network under rip.
`(config-router-af)#exit`	Exit address-family mode.
`(config-router)#ex`	Exit router mode
`(config)#interface eth1`	Enter interface mode

(config-if)#ip vrf forwarding vrf1	Associate the interface to vrf1
(config-if)#ip address 2.2.2.1/24	Configure the IP address 2.2.2.1 to eth1

R2

#configure terminal	Enter configure mode.
(config)#virtual-router VR1	Create virtual router VR1.
(config-vr)#load rip	Load the rip module in VR1.
(config)#exit	Exit VR mode.
(config)#interface eth1	Enter interface mode.
(config-if)#virtual-router forwarding VR1	Associate eth1 to VR1.
(config-if)#exit	Exit interface mode.
(config)#exit	Exit configure mode.
#login virtual-router VR1	Log in to virtual router VR1.
>en	Enter privileged exec mode.
#con t	Enter configuration mode.
(config)#ip vrf vrf1	Create vrf1
((config-vrf)#exit	Exit VRF mode
(config)#router rip	Enter router mode.
(config-router)#version 2	Specify RIP version 2.
(config-router)#address-family ipv4 vrf vrf1	Enter address family mode for vrf1
(config-router-af)#network 2.2.2.0/24	Advertise the connected network under RIP.
(config-router-af)#exit	Exit address family mode.
(config-router)#ex	Exit router mode
(config)#interface eth1	Enter interface mode
(config-if)#ip vrf forwarding vrf1	Associate the interface to vrf1
(config-if)#ip address 2.2.2.2/24	Configure the IP address 2.2.2.1 to eth1

Validation

Verify the routing table in R1:

```
#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
       area
       * - candidate default
IP Route Table for VRF "vrf1"
C       2.2.2.0/24 is directly connected, eth1
Gateway of last resort is not set
```

Verify RIP database in R1:

```
#show ip rip database vrf vrf1
Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
```

```
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP
    Network            Next Hop         Metric  From            If     Time
   Rc 2.2.2.0/24                              1                        eth1
```

Verify the routing table in R2:

```
#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
       area
       * - candidate default
IP Route Table for VRF "vrf1"
C       2.2.2.0/24 is directly connected, eth1
Gateway of last resort is not set
```

Verify RIP database in R2:

```
#show ip rip database vrf vrf1
Codes: R - RIP, Rc - RIP connected, Rs - RIP static, K - Kernel,
       C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP
   Network            Next Hop         Metric From            If     Time
  Rc 2.2.2.0/24                              1                        eth1
```

Configuration IPv6 VRF

R1

`#configure terminal`	Enter the Configure mode.
`(config)#virtual-router VR1`	Create a virtual router VR1.
`(config-vr)#load ipv6 rip`	Load the RIPng module to be used in VR1.
`(config)#exit`	Exit the VR mode.
`(config)#interface eth1`	Switch to interface eth1.
`(config-if)#virtual-router forwarding VR1`	Associate eth1 to VR1.
`(config-if)#exit`	Exit interface mode.
`(config)#exit`	Exit the router config mode.
`#login virtual-router VR1`	Login in to virtual-router VR1.
`>en`	Enable the config mode.
`#con t`	Enter the router configuration mode.
`(config)#ip vrf vrf1`	Create the vrf1
`(config-vrf)#exit`	Exit the vrf
`(config)#router ipv6 rip`	Enable ipv6 rip
`(config-router)#address-family ipv6 vrf vrf1`	Switch to ipv6 rip address family
`(config-router-af)#aggregate-address 2222::/48`	Configure the ipv6 aggregate-address
`(config-router-af)#exit`	Exit the router rip mode.
`(config-router)#ex`	exit.

(config)#interface eth1	Switch to interface eth1
(config-if)#ip vrf forwarding vrf1	Associate eth1 to the vrf1
(config-if)#ip address 2.2.2.1/24	Configure the ip address 2.2.2.2 to interface eth1
(config-if)#ipv6 address 2222::1/48	Configure the ipv6 address.
(config-if)#ipv6 address fe80::1/48	Configure the link local address
(config-if)#ipv6 router rip	Associate interface eth1 to ipv6 rip

R2

#configure terminal	Enter the Configure mode.
(config)#virtual-router VR1	Create a virtual router VR1.
(config-vr)#load ipv6 rip	Load the Ripng module to be used in VR1
(config)#exit	Exit the VR mode.
(config)#interface eth1	Switch to interface eth1.
(config-if)#virtual-router forwarding VR1	Associate eth1 to VR1.
(config-if)#exit	Exit interface mode.
(config)#exit	Exit the router config mode.
#login virtual-router VR1	Login in to virtual-router VR1.
>en	Enable the config mode.
#con t	Enter the router configuration mode.
(config)#ip vrf vrf1	Create the vrf1
(config-vrf)#exit	Exit the vrf
(config)#router ipv6 rip	Switch to the rip ipv6 address family
(config-router)#address-family ipv6 vrf vrf1	Switch to the rip ipv6 address family
(config-router-af)#aggregate-address 2222::/48	Configure the ipv6 aggregate address
(config-router-af)#exit	Exit the address ipv6 rip address family
(config-router)#ex	Exit the router IPV6 rip mode.
(config)#interface eth1	Exit to the router config mode
(config-if)#ip vrf forwarding vrf	Associate eth1 to the vrf1
(config-if)#ip address 2.2.2.2/24	Configure the ip address 2.2.2.1 to eth eth1
(config-if)#ipv6 address 2222::2/48	Configure the ipv6 address.
(config-if)#ipv6 address fe80::2/48	Configure the link local address
(config-if)#ipv6 router rip	Associate interface eth1 to ipv6 rip

Validation

```
rtr1#show ipv6 route vrf vrf1
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, I - IS-IS, B - BGP
Timers: Uptime
IP Route Table for VRF "vrf1"
C      2222::/48 via ::, eth1, 00:06:19
```

```
C      fe80::/48 via ::, eth1, 00:02:33
rtr18#show ipv6 rip database vrf vrf1
Codes: R - RIP, Rc - RIP connected, Rs - RIP static, Ra - RIP aggregated,
       Rcx - RIP connect suppressed, Rsx - RIP static suppressed,
       K - Kernel, C - Connected, S - Static, O - OSPF, I - IS-IS, B - BGP
   Network                    Next Hop                  If     Met Tag  Time
Rcx 2222::/48                 ::                        eth1     1   0

rtr2#show ipv6 rip interface
eth1 is up, line protocol is up
  Routing Protocol: RIPng
    VPN Routing/Forwarding: vrf1
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IPv6 interface address:
      2222::1/48
      fe80::1/48
eth2 is up, line protocol is up
  RIPng is not enabled on this interface
```

CHAPTER 4 ISIS Configuration

Overview

Intermediate System to Intermediate System (IS-IS) is an interior routing protocol operating within a single administrative domain.It is a link-state routing protocol, operating by reliably flooding link state information throughout a network of routers. Each IS-IS router independently builds a database of the network's topology, aggregating the flooded network information. IS-IS uses Dijkstra's algorithm for computing the best path through the network. Packets (datagrams) are then forwarded, based on the computed ideal path, through the network to the destination.

Topology

Figure 4-71: ISIS Topology for VRF

Configuration IPv4 VRF

R1

`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#exit`	Exit VRF mode
`(config)#router isis 1 vrf1`	Associate the ISIS process to vrf1
`(config-router)#net 49.0001.0000.0000.0001.00`	Establish a network entity title for this instance, specifying the area address and the system ID.
`(config-router)#is-type level-1`	Configure instance 1 as level-1 routing
`(config-router)#exit`	Exit router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf1`	Associate the interface to vrf1
`(config-if)#ip address 2.2.2.1/24`	Configure the IP address 2.2.2.1 to eth1
`(config-if)#ip router isis 1`	Enable ISIS routing on an instance for area 49
`(config-if)#isis circuit-type level-1`	Configure interface as level-1
`(config-if)#exit`	Exit interface mode
`(config)#exit`	Exit config mode

R2

`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#exit`	Exit VRF mode
`(config)#router isis 1 vrf1`	Associate the ISIS process to vrf1
`(config-router)#net 49.0001.0000.0000.0002.00`	Establish a network entity title for this instance, specifying the area address and the system ID.
`(config-router)#is-type level-1`	Configure instance 1 as level-1 routing
`(config-router)#exit`	Exit router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf1`	Associate the interface to vrf1
`(config-if)#ip address 2.2.2.2/24`	Configure the IP address 2.2.2.1 to eth1
`(config-if)#ip router isis 1`	Enable ISIS routing on an instance for area 49
`(config-if)#isis circuit-type level-1`	Configure interface as level-1
`(config-if)#exit`	Exit interface mode
`(config)#exit`	Exit config mode

Validation

R1

```
#show clns neighbors
Tag 1:  VRF : vrf1
System Id      Interface   SNPA      State     Holdtime        Type     Protocol
0000.0000.0002      eth1   5254.00d1.5789      Up          9    L1      IS-IS
```

R2

```
#show clns neighbors
Tag 1:  VRF : vrf1
System Id      Interface   SNPA                State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.0024.4323      Up     26        L1   IS-IS
```

R1

```
#show ip isis route
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric

Tag 1:  VRF : vrf1
     Destination         Metric      Next-Hop            Interface      Tag
C    2.2.2.0/24         10         --                eth1          0            0
```

R2

```
#show ip isis route
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
```

```
Tag 1:  VRF : vrf1
     Destination         Metric      Next-Hop            Interface      Tag
C    2.2.2.0/24         10         --                eth1          0            0
```

CHAPTER 5 BGP Configuration

Overview

Border Gateway Protocol (BGP) makes core routing decisions on the Internet using a table of IP networks or “prefixes” which designate network reachability among autonomous systems (AS). BGP is a path vector protocol or a variant of a distance-vector routing protocol. BGP does not involve traditional Interior Gateway Protocol (IGP) metrics, but routing decisions are made based on path, network policies, and/or rule sets. For this reason, it is more appropriately termed a reachability protocol rather than routing protocol.

Note: This chapter covers BGP configuration in non-default VR and non-default VRF.

Enabling dynamic capability on the specific BGP peer will prevent BGP session to reset.

Topology

Figure 5-72: BGP topology for VRF

Configuration

R1

`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#rd 800:1`	Specify the route distinguisher in the VRF
`(config-vrf)#route-target import 800:1`	Specify the import route target
`(config-vrf)#route-target export 800:1`	Specify the export route target
`(config-vrf)#exit`	Exit VRF mode
`(config)#router bgp 200`	Enter the bgp configuration mode
`(config-router)#address-family ipv4 vrf  vrf1`	Enter address family mode for vrf1
`(config-router-af)#neighbor 2.2.2.2 remote-as 100`	Specify the BGP neighbor and remote-AS.
`(config-router-af)#exit`	Exit address family mode.
`(config-router)#ex`	Exit router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf1`	Associate the interface to vrf1
`(config-if)#ip address 2.2.2.1/24`	Configure the IP address 2.2.2.1 to eth1

R2

`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#rd 800:1`	Specify the route distinguisher in the VRF
`(config-vrf)#route-target import 800:1`	Specify the import route target
`(config-vrf)#route-target export 800:1`	Specify the export route target
`(config-vrf)#exit`	Exit vrf mode
`(config)#router bgp 100`	Enter router mode.
`(config-router)#address-family ipv4 vrf  vrf1`	Enter address family mode for vrf1
`(config-router-af)#neighbour 2.2.2.1 remote-as 200`	Specify the BGP neighbor and remote-as.
`(config-router-af)#exit`	Exit address family mode.
`(config-router)#ex`	Exit router mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding vrf1`	Associate the interface to vrf1
`(config-if)#ip address 2.2.2.2/24`	Configure the IP address 2.2.2.1 to eth1

Validation

Verify the routing table in R1

```
#show ip bgp neighbors
BGP neighbor is 2.2.2.2, vrf vrf1, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 2.2.2.1, remote router ID 2.2.2.2
  BGP state = Established, up for 00:00:14
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 2 messages, 0 notifications, 0 in queue
  Sent 3 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 2.2.2.1, Local port: 179
Foreign host: 2.2.2.2, Foreign port: 36200
Nexthop: 2.2.2.1
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network
```

Verify the routing table in R2

```
#show ip bgp neighbors
BBGP neighbor is 2.2.2.1, vrf vrf1, remote AS 200, local AS 100, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 2.2.2.1
  BGP state = Established, up for 00:08:09
  Last read 00:00:09, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 18 messages, 0 notifications, 0 in queue
  Sent 18 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 2.2.2.2, Local port: 36200
Foreign host: 2.2.2.1, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

R1

```
#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "vrf1"
C     2.2.2.0/24 is directly connected, eth1, 00:20:40

Gateway of last resort is not set
```

R2

```
#show ip route vrf vrf1
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "vrf1"
C     2.2.2.0/24 is directly connected, eth1, 00:14:53

Gateway of last resort is not set
```

CHAPTER 6 Inter-VRF Route Leaking Configuration

This chapter shows how to configure inter-VRF route leaking.

Overview

Virtual Routing and Forwarding (VRF) provides the ability to have multiple virtual routers on a single physical device. VRFs operate without knowledge of one another unless they are imported or exported into one another using inter-VRF route leaking. Inter-VRF route leaking allows leaking of route prefixes from one VRF instance to another VRF instance on the same physical router which eliminates the need for external routing.This is useful in cases where multiple VRFs share the same path to reach an external domain, while maintaining their internal routing information limited to their own VRF. This feature enables a data center to consolidate multiple VRF services onto a single server.

There are two types of inter-VRF route leaking:

- Static leaking: leaking manually configured static route entries from a source VRF to a global default VRF table.
- Dynamic leaking: leaking connected routes and dynamically learned routes from protocols such as ISIS, OSPF, and BGP from a source VRF to a destination VRF.

Static Leaking

Static route leaking directly between VRFs is not supported. What does work is routing traffic from a VRF to the global default VRF routing table. One advantage of using static route leaking is that you can configure exactly which routes are reachable without configuring BGP.

Topology

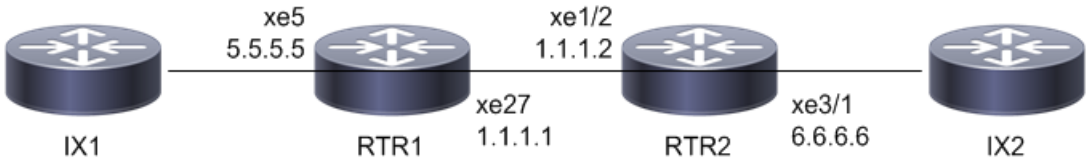

Figure 6-73: Static leaking

Configuration

The following steps describe how to configure static leaking.

RTR1

`RTR1#configure terminal`	Enter configure mode.
`RTR1(config)#interface xe5`	Enter interface mode
`RTR1(config-if)#ip address 5.5.5.5/24`	Assign IP address 5.5.5.5 to interface xe5
`RTR1(config-if)#exit`	Exit interface mode
`RTR1(config)#interface xe27`	Enter interface mode

`RTR1(config-if)#ip address 1.1.1.1/24`	Assign IP address 1.1.1.1 to interface xe27
`RTR1(config-if)#exit`	Exit interface mode

RTR2

`RTR2#configure terminal`	Enter configure mode.
`RTR2(config)#interface xe1/2`	Enter interface mode
`RTR2(config-if)#ip address 1.1.1.2/24`	Assign IP address 1.1.1.2 to interface xe1/2
`RTR2(config-if)#exit`	Exit interface mode
`RTR2(config)#ip vrf vrf1`	Create VRF vrf1
`RTR2(config-vrf)#exit`	Exit VRF mode
`RTR2(config)#interface xe3/1`	Enter interface mode
`RTR2(config-if)#ip vrf forwarding vrf1`	Associate xe3/1 to vrf1
`RTR2(config-if)#ip address 6.6.6.6/24`	Assign IP address 6.6.6.6 to interface xe3/1
`RTR2(config-if)#exit`	Exit interface mode
`RTR2(config)#ip route vrf vrf1 5.5.5.0/24 1.1.1.1 xe1/2 global`	Add static route to reach global default VRF table
`RTR2(config)#exit`	Exit configure mode

Validation

```
RTR2#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     1.1.1.0/24 is directly connected, xe1/2, 00:00:05
C     127.0.0.0/8 is directly connected, lo, 00:20:38
IP Route Table for VRF "management"
Gateway of last resort is 10.12.29.1 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 10.12.29.1, eth0, 00:20:38
C     10.12.29.0/24 is directly connected, eth0, 00:20:38
C     127.0.0.0/8 is directly connected, lo.management, 00:20:38
IP Route Table for VRF "vrf1"
S    v5.5.5.0/24 [1/0] via 1.1.1.1, xe1/2, 00:05:20
C     6.6.6.0/24 is directly connected, xe3/1, 00:07:06
C     127.0.0.0/8 is directly connected, lo.vrf1, 00:12:25

RTR2#show ip route vrf all database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  1.1.1.0/24 is directly connected, xe1/2, 00:00:51
C    *>  127.0.0.0/8 is directly connected, lo, 00:21:24
IP Route Table for VRF "management"
S    *>  0.0.0.0/0 [1/0] via 10.12.29.1, eth0, 00:21:24
C    *>  10.12.29.0/24 is directly connected, eth0, 00:21:24
C    *>  127.0.0.0/8 is directly connected, lo.management, 00:21:24
IP Route Table for VRF "vrf1"
S    *> v5.5.5.0/24 [1/0] via 1.1.1.1, xe1/2, 00:06:06
C    *>  6.6.6.0/24 is directly connected, xe3/1, 00:07:52
C    *>  127.0.0.0/8 is directly connected, lo.vrf1, 00:13:11
```

Dynamic Leaking

Route Leaking enables communication between isolated (virtual) routing domains by segregating and sharing a set of services that are available on one routing domain with other virtual domains. Inter-VRF route leaking enables a VRF to leak or export routes in its router to one or more VRFs. Dynamic route leaking enables a source VRF to share both its connected routes as well as dynamically learned routes from protocols such as ISIS, OSPF, and BGP to destination VRFs.

Topology

Figure 6-74: Dynamic leaking

Configuration

The following steps describe how to configure dynamic leaking.

RTR1

`RTR1#configure terminal`	Enter configure mode.
`RTR1(config)#interface xe5`	Enter interface mode
`RTR1(config-if)#ip address 5.5.5.5/24`	Assign IP address 5.5.5.5 to interface xe5
`RTR1(config-if)#exit`	Exit interface mode
`RTR1(config)#interface xe27`	Enter interface mode
`RTR1(config-if)#ip address 1.1.1.1/24`	Assign IP address 1.1.1.1 to interface xe27

`RTR1(config-if)#exit`	Exit interface mode
`RTR1(config)#router ospf 1`	Enter OSPF router mode
`RTR1(config-router)#network 1.1.1.0/24 area 0`	Specify the network type and area 0
`RTR1(config-router)#redistribute connected`	Redistribute connected route inside ospf
`RTR1(config-router)#exit`	Exit OSPF router mode

RTR2

`RTR2#configure terminal`	Enter configure mode.
`RTR2(config)#ip vrf vrf1`	Create VRF vrf1
`RTR2(config-vrf)#rd 100:1`	Configure route distinguisher in the VRF
`RTR2(config-vrf)#route-target export 100:1`	Configure export route target
`RTR2(config-vrf)#route-target import 200:1`	Configure import route target
`RTR2(config-vrf)#exit`	Exit VRF mode
`RTR2(config)#ip vrf vrf2`	Create VRF vrf2
`RTR2(config-vrf)#rd 200:1`	Configure route distinguisher in the VRF
`RTR2(config-vrf)#route-target export 200:1`	Configure export route target
`RTR2(config-vrf)#route-target import 100:1`	Configure import route target
`RTR2(config-vrf)#exit`	Exit VRF mode
`RTR2(config)#interface xe1/2`	Enter interface mode
`RTR2(config-if)#ip vrf forwarding vrf1`	Associate vrf1 to interface xe1/2
`RTR2(config-if)#ip address 1.1.1.2/24`	Assign IP address 1.1.1.2 to interface xe1/2
`RTR2(config-if)#exit`	Exit interface mode
`RTR2(config)#interface xe3/3`	Enter interface mode
`RTR2(config-if)#ip vrf forwarding vrf2`	Associate vrf2 to interface xe3/3
`RTR2(config-if)#ip address 2.2.2.2/24`	Assign IP address 2.2.2.2 to interface xe3/3
`RTR2(config-if)#exit`	Exit interface mode
`RTR2(config)#router ospf 1 vrf1`	Associate the OSPF process with vrf1
`RTR2(config-router)#network 1.1.1.0/24 area 0`	Specify the network type and area 0
`RTR2(config-router)#redistribute bgp`	Redistribute BGP routes inside OSPF
`RTR2(config-router)#exit`	Exit router mode
`RTR2(config)#router ospf 2 vrf2`	Associate the OSPF process with vrf2
`RTR2(config-router)#network 2.2.2.0/24 area 0`	Specify the network type and area 0
`RTR2(config-router)#redistribute bgp`	Redistribute BGP routes inside OSPF
`RTR2(config-router)#exit`	Exit router mode
`RTR2(config)#router bgp 100`	Enter BGP router mode
`RTR2(config-router)#address-family ipv4 vrf vrf1`	Enter address family mode for vrf1
`RTR2(config-router-af)#redistribute ospf 1`	Redistribute OSPF routes inside BGP
`RTR2(config-router-af)#exit-address-family`	Exit address family mode

`RTR2(config-router)#address-family ipv4 vrf vrf2`	Enter address family mode for vrf2
`RTR2(config-router-af)#redistribute ospf 2`	Redistribute OSPF routes inside BGP
`RTR2(config-router-af)#exit-address-family`	Exit address family mode
`RTR2(config-router)#exit`	Exit router mode

RTR3

`RTR3#configure terminal`	Enter configure mode.
`RTR3(config)#interface xe1`	Enter interface mode
`RTR3(config-if)#ip address 6.6.6.6/24`	Assign IP address 6.6.6.6 to interface xe1
`RTR3(config-if)#exit`	Exit from config mode
`RTR3(config)#interface xe33`	Enter interface mode
`RTR3(config-if)#ip address 2.2.2.3/24`	Assign IP address 2.2.2.3 to interface xe33
`RTR3(config-if)#exit`	Exit interface mode
`RTR3(config)#router ospf 2`	Enter OSPF router mode
`RTR3(config-router)#network 2.2.2.0/24 area 0`	Specify the network type and area 0
`RTR3(config-router)#redistribute connected`	Redistribute connected route inside ospf
`RTR3(config-router)#exit`	Exit OSPF router mode

Validation

RTR1

```
RTR1#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
1.1.1.2           1   Full/Backup     00:00:39    1.1.1.2         xe27

RTR1#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     1.1.1.0/24 is directly connected, xe27, 01:51:47
O E2  2.2.2.0/24 [110/1] via 1.1.1.2, xe27, 00:22:51
C     5.5.5.0/24 is directly connected, xe5, 02:16:39
O E2  6.6.6.0/24 [110/1] via 1.1.1.2, xe27, 00:22:51
```

```
C     127.0.0.0/8 is directly connected, lo, 02:25:23
```

RTR2

```
RTR2#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(vrf1):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
5.5.5.5           1   Full/DR         00:00:34    1.1.1.1         xe1/2                  0

Total number of full neighbors: 1
OSPF process 2 VRF(vrf2):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
6.6.6.6           1   Full/DR         00:00:36    2.2.2.3         xe3/3                  0

RTR2#sh ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     127.0.0.0/8 is directly connected, lo, 02:06:03
IP Route Table for VRF "management"
Gateway of last resort is 10.12.29.1 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 10.12.29.1, eth0, 02:06:03
C     10.12.29.0/24 is directly connected, eth0, 02:06:03
C     127.0.0.0/8 is directly connected, lo.management, 02:06:03
IP Route Table for VRF "vrf1"
C     1.1.1.0/24 is directly connected, xe1/2, 01:31:20
B    v2.2.2.0/24 [20/1] is directly connected, xe3/3, 00:02:35
O E2  5.5.5.0/24 [110/20] via 1.1.1.1, xe1/2, 00:07:12
B    v6.6.6.0/24 [20/20] via 2.2.2.3, xe3/3, 00:02:35
C     127.0.0.0/8 is directly connected, lo.vrf1, 01:40:49
IP Route Table for VRF "vrf2"
B    v1.1.1.0/24 [20/1] is directly connected, xe1/2, 00:03:35
C     2.2.2.0/24 is directly connected, xe3/3, 01:31:02
B    v5.5.5.0/24 [20/20] via 1.1.1.1, xe1/2, 00:03:35
O E2  6.6.6.0/24 [110/20] via 2.2.2.3, xe3/3, 00:06:52
C     127.0.0.0/8 is directly connected, lo.vrf2, 01:32:22

RTR2#sh ip route vrf all database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *>  127.0.0.0/8 is directly connected, lo, 02:07:34
IP Route Table for VRF "management"
S    *>  0.0.0.0/0 [1/0] via 10.12.29.1, eth0, 02:07:34
C    *>  10.12.29.0/24 is directly connected, eth0, 02:07:34
C    *>  127.0.0.0/8 is directly connected, lo.management, 02:07:34
IP Route Table for VRF "vrf1"
C    *>  1.1.1.0/24 is directly connected, xe1/2, 01:32:51
O        1.1.1.0/24 [110/1] is directly connected, xe1/2, 00:09:13
B    *> v2.2.2.0/24 [20/1] is directly connected, xe3/3, 00:04:06
O E2 *>  5.5.5.0/24 [110/20] via 1.1.1.1, xe1/2, 00:08:43
B    *> v6.6.6.0/24 [20/20] via 2.2.2.3, xe3/3, 00:04:06
C    *>  127.0.0.0/8 is directly connected, lo.vrf1, 01:42:20
IP Route Table for VRF "vrf2"
B    *> v1.1.1.0/24 [20/1] is directly connected, xe1/2, 00:05:06
C    *>  2.2.2.0/24 is directly connected, xe3/3, 01:32:33
O        2.2.2.0/24 [110/1] is directly connected, xe3/3, 00:08:42
B    *> v5.5.5.0/24 [20/20] via 1.1.1.1, xe1/2, 00:05:06
O E2 *>  6.6.6.0/24 [110/20] via 2.2.2.3, xe3/3, 00:08:23
C    *>  127.0.0.0/8 is directly connected, lo.vrf2, 01:33:53
```

RTR3

```
RTR3#sh ip ospf neighbor

Total number of full neighbors: 1
OSPF process 2 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
2.2.2.2           1   Full/Backup     00:00:37    2.2.2.2         xe33                    0

RTR3#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       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
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O E2  1.1.1.0/24 [110/1] via 2.2.2.2, xe33, 00:20:12
C     2.2.2.0/24 is directly connected, xe33, 01:47:45
O E2  5.5.5.0/24 [110/1] via 2.2.2.2, xe33, 00:20:12
C     6.6.6.0/24 is directly connected, xe1, 02:00:13
C     127.0.0.0/8 is directly connected, lo, 02:21:14
```

Virtual Router Redundancy Protocol Configuration

CHAPTER 1 VRRP Configuration

This chapter provides an overview of Virtual Router Redundancy Protocol (VRRP) and its implementation with OcNOS. VRRP eliminates the risk of a single point of failure inherent in a static default routing environment. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. One of the major advantages of VRRP is that it makes default path available without requiring configuration of dynamic routing on every end-host.

OcNOS only supports VRRP protocol version 3.

Terminology

Terms related to VRRP configuration are defined in the table below.

Backup Router	The VRRP router that is backing up an IP address. It assumes forwarding responsibility for the virtual IP address if the Master fails.
Critical IP	The IP address that a VRRP router sends/receives messages on for a particular session.
IP Address Owner	The VRRP Router that has the virtual router's IP address(es) as real interface address(es). This is the router that, when up, will respond to packets addressed to one of these IP addresses for ICMP pings, TCP connections, and so on
Master Router	The VRRP router that owns the IP address (i.e., is being backed up), and which is the default router for forwarding for that IP address.
Virtual IP	The IP address that is being backed up by a VRRP session.
Virtual Router	A router managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier and a set of associated IP addresses across a common LAN. A VRRP Router might backup one or more virtual routers.
VRRPv2 Router	A router running the Virtual Router Redundancy Protocol version 2. It might participate in one or more virtual routers.
VRRPv3 Router	A router running the Virtual Router Redundancy Protocol version 3. It might participate in one or more virtual routers.

VRRP Process

Typically, end hosts are connected to the enterprise network through a single router (first-hop router) that is in the same Local Area Network (LAN) segment. The most popular method of configuration is for the end hosts to configure statically this router as their default gateway. This minimizes configuration and processing overhead. As shown in Figure 1-75, the problem with this configuration is that it produces a single point of failure if this first-hop router fails.

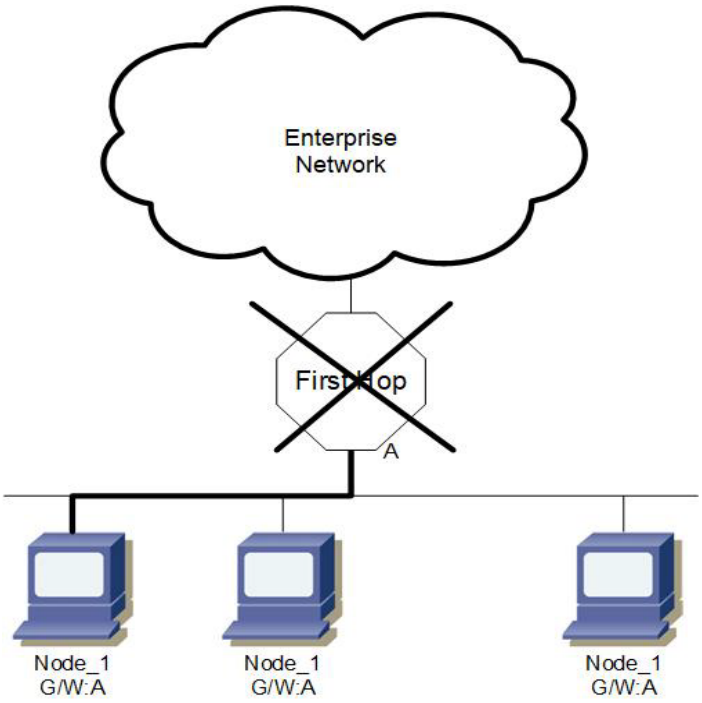

Figure 1-75: VRRP Process - First-Hop Not Reachable

The Virtual Router Redundancy Protocol attempts to solve this problem by introducing the concept of a virtual router, composed of two or more VRRP routers on the same subnet as shown in Figure 1-76. The concept of a virtual IP address is also introduced, which is the address that end hosts configure as their default gateway. One of the routers called the “Master” forwards packets on behalf of this IP address.

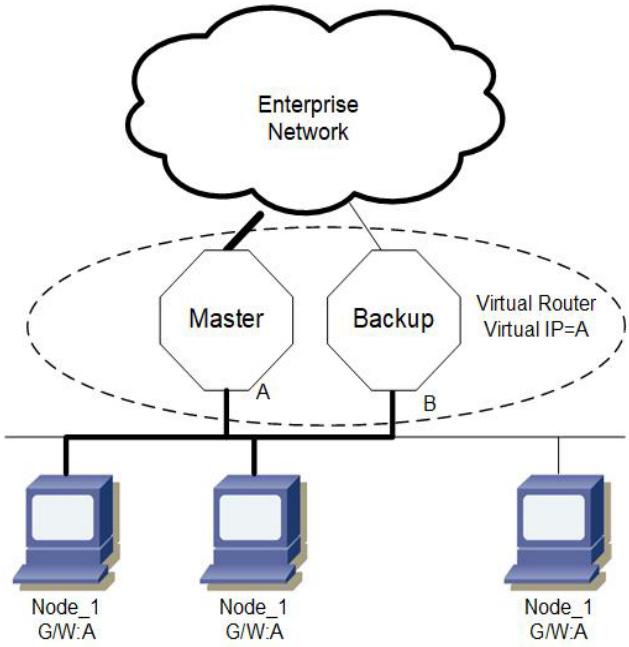

Figure 1-76: VRRP Process - Master and Backup VR

As shown in Figure 1-77, if the Master router fails, one of the other routers (Backup) assumes forwarding responsibility for it.

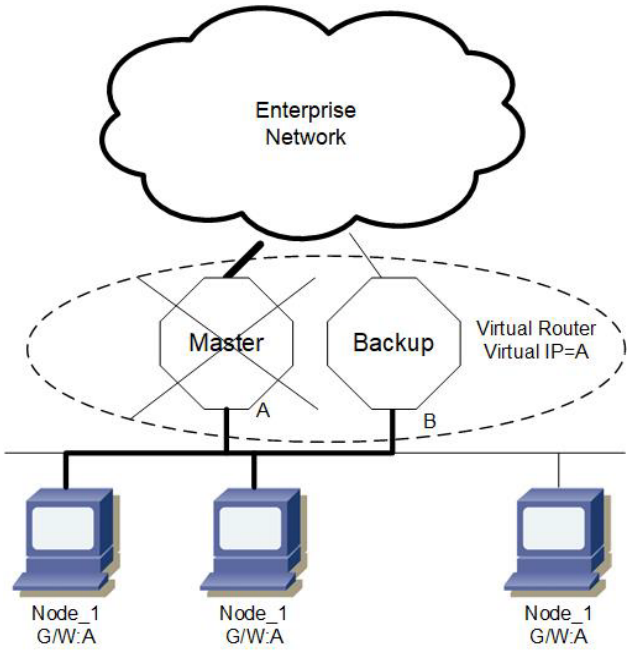

Figure 1-77: VRRP Process - Master Down and Backup Takeover

At first glance, the configuration in might not seem very useful, as it doubles the cost, and leaves one router idle at all times. This, however, can be avoided by creating two virtual routers and splitting the traffic between them.

Note: Adding a default route in the kernel on the interface that is used for VRRP might cause loss of network connectivity. According to the VRRP guidelines, when the VRRP session changes, the MAC address for the machine that attains the master state also changes. The change causes the default route from the kernel to disappear and leads to loss of connectivity. To avoid this situation, add the default route in the NSM and not in the kernel. This ensures that the default route remains on the machine across changes in the VRRP state.

To add default route through NSM, run the following command in NSM:

```
ip route 0.0.0.0/0 <IPADDRESS>
```

where `<IPADDRESS>` is the IP address of the default gateway.

One Virtual Router

In this configuration, the end-hosts install a default route to the IP address of virtual router 1(VRID = 1), and both routers R1 and R2 run VRRP. R1 is configured to be the Owner for virtual router 1 (VRID = 1) and R2 as a Backup for virtual router 1. If R1 fails, R2 will take over virtual router 1 and its IP addresses, and provide uninterrupted service for the hosts. Configuring only one virtual router doubles the cost, and leaves R2 idle at all times.

Topology

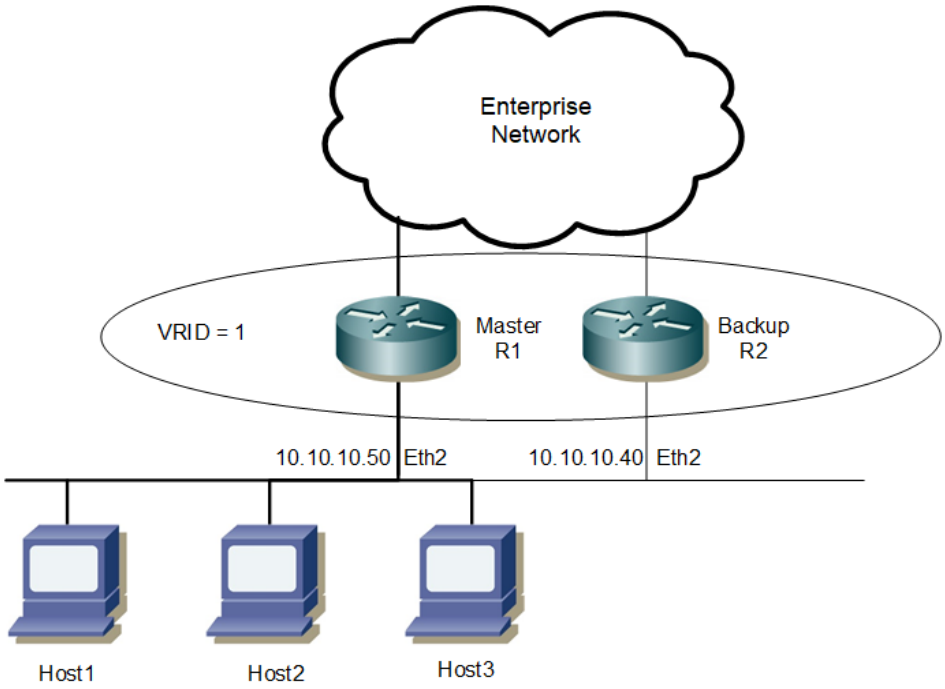

Figure 1-78: VRRP with One Virtual Router

R1

`#configure terminal`	Enter the `Configure` mode.
`(config)#router vrrp 1 eth2`	Create a VRRP instance for interface eth2.
`(config-router)#virtual-ip 10.10.10.50 owner`	Set the virtual IP address for the VRRP session. Define the default state (owner) of the VRRP router within the virtual router.
`(config-router)#preempt-mode true`	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
`(config-router)#advertisement-interval 100`	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
`(config-router)# v2-compatible`	Enable the v2-compatible
`(config-router)#authentication text abcd`	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
`(config-router)#enable`	Enable the VRRP session on the router.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode.

R2

`#configure terminal`	Enter the `Configure` mode.
`(config)#router vrrp 1 eth2`	Create a VRRP instance for interface eth2.
`(config-router)#virtual-ip 10.10.10.50`	Set the virtual IP address for the VRRP session.

(config-router)#priority 200	Configure the priority to 200 (less than 255), because R2 is the Backup router.
(config-router)#preempt-mode true	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
(config-router)#advertisement-interval 100	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

Validation

DUT

```
#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility enabled

Address family IPv4
 VRRP Id: 1 on interface: vlan1.1000
 State: AdminUp   - Master
 Virtual IP address: 10.1.1.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.1.1.2
 Operational master IP address: 10.1.1.2
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 1 minutes 52 seconds (11200 centi sec)
 Master uptime: 0 hours 0 minutes 21 seconds (2100 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv4 interface vlan1.1000: JOINED
 V2-Compatible: TRUE
```

SD-1

```
#show vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminUp   - Backup
 Virtual IP address: 10.10.10.1 (Not-owner)
```

```
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.10.10.40
Operational master IP address: 10.10.10.50
Priority is 90
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 0 minutes 29 seconds (2900 centi sec)
Skew time: 80 centi sec
Master Down Interval: 380 centi sec
Accept mode: TRUE
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface eth2: JOINED
V2-Compatible: TRUE
```

Two Virtual Routers

In the previous, one virtual router example, R2 is not backed up by R1. This example illustrates how to back up R2 by configuring a second virtual router. In this configuration, R1 and R2 are two virtual routers, and the hosts split their traffic between R1 and R2. R1and R2 functions as backups for each other.

Topology

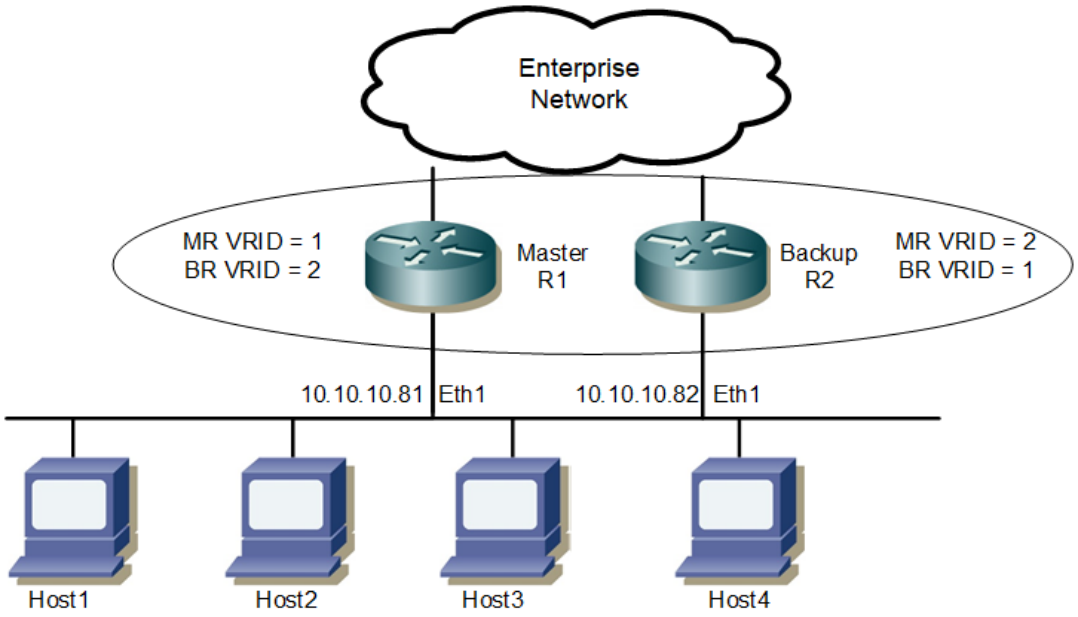

Figure 1-79: Configuring VRRP with Two Virtual Routers

R1

`#configure terminal`	Enter the `Configure` mode.
`(config)#router vrrp 1 xe1`	Create a VRRP instance for interface xe1.

(config-router)#virtual-ip 10.10.10.81 owner	Set the virtual IP address for the VRRP session. Define the default state (owner) of the VRRP router within the virtual router.
(config-router)#preempt-mode true	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
(config-router)#advertisement-interval 100	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session 1 on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode and enter the Configure mode.
(config)#router vrrp 2 xe1	Create a VRRP instance for interface xe1.
(config-router)#virtual-ip 10.10.10.82	Set the virtual IP address for the VRRP session.
(config-router)#priority 200	Configure the priority to 200 (less than 255), because R2 is the Backup router.
(config-router)#preempt-mode true	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
(config-router)#advertisement-interval 100	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session two on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

R2

#configure terminal	Enter the Configure mode.
(config)#router vrrp 1 xe1	Create a VRRP instance for interface xe1.
(config-router)#virtual-ip 10.10.10.81	Set the virtual IP address for the VRRP session.
(config-router)#priority 200	Configure the priority to 200 (less than 255), because R2 is the Backup router.
(config-router)#preempt-mode true	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
(config-router)#advertisement-interval 100	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session 1 on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.

(config-router)#exit	Exit the Router mode and enter the Configure mode.
(config)#router vrrp 2 xe1	Create a VRRP instance for interface xe1
(config-router)#virtual-ip 10.10.10.82 owner	Set the virtual IP address for the VRRP session. Define the default state (owner) of the VRRP router within the virtual router.
(config-router)#preempt-mode true	Set the preempt mode to specify that the highest priority will function as a backup to master when master is unavailable.
(config-router)#advertisement-interval 100	Configure the advertisement interval to 100 centi seconds (value must be a multiple of 5).
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session two on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

Validation

The following outputs on R1and R2 display the complete configuration for each session on R1 and R2. In session one, R1 is the master router, and in session two R2 is the master router.

R1

```
R1#sh vrrp 1 eth1 VRRP Version: 3 VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: xe1
State: AdminUp- Master
Virtual IP address: 10.10.10.81 (Owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.10.10.81
Operational master IP address: 10.10.10.81
Priority is 255
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 1 minutes 25 seconds (8500 centi sec)
Master uptime: 0 hours 1 minutes 18 seconds (7800 centi sec)
Accept mode: FALSE Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface eth1: JOINED V2-Compatible: TRUE

R1#sh vrrp 2 eth1 VRRP Version: 3 VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 2 on interface: xe1
State: AdminUp- Backup
Virtual IP address: 10.10.10.82 (Not-owner)
```

```
Virtual MAC address is 0000.5e00.0102
Operational primary IP address: 10.10.10.81
Operational master IP address: 10.10.10.82 Priority is 200
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 2 minutes 3 seconds (12300 centi sec)
Skew time: 80 centi sec
Master Down Interval: 380 centi sec
Accept mode: FALSE Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface eth1: JOINED V2-Compatible: TRUE
```

R2

```
R2#sh vrrp 1 eth1 VRRP Version: 3 VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: xe1
State: AdminUp- Backup
Virtual IP address: 10.10.10.81 (Not-owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.10.10.82
Operational master IP address: 10.10.10.81
Priority is 200
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 0 minutes 37 seconds (3700 centi sec)
Skew time: 80 centi sec
Master Down Interval: 380 centi sec
Accept mode: FALSE Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface eth1: JOINED
V2-Compatible: TRUE

R2#sh vrrp 2 eth1 VRRP Version: 3 VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 2 on interface: xe1
State: AdminUp- Master
Virtual IP address: 10.10.10.82 (Owner)
Virtual MAC address is 0000.5e00.0102
Operational primary IP address: 10.10.10.82
Operational master IP address: 10.10.10.82 Priority is 255
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 0 minutes 12 seconds (1200 centi sec)
Master uptime: 0 hours 1 minutes 18 seconds (7800 centi sec)
Accept mode: FALSE
```

```
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface eth1: JOINED
V2-Compatible: TRUE

R2#
```

Two Backup Routers

In this configuration, Host B could be a gateway router. As such, interface eth1 on Routers R1, R2, and R3, and the gateway router, would run the IGP protocol.

Topology

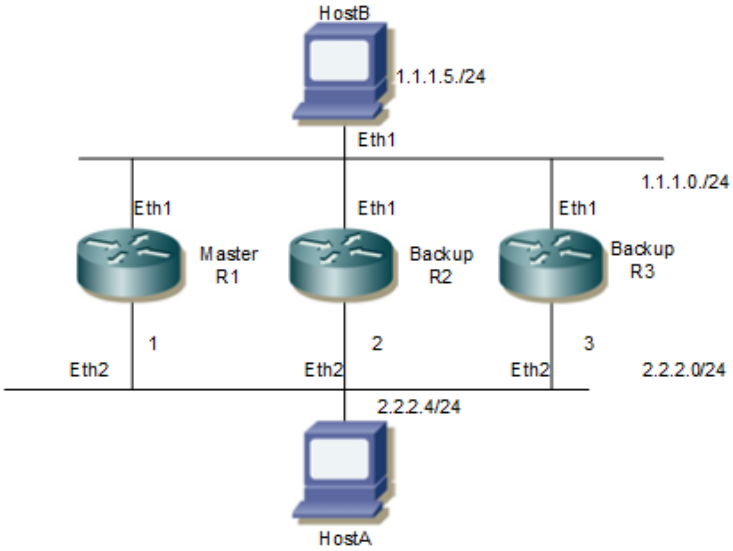

Figure 1-80: Configuring VRRP with Two Backup Routers

R1

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth2`	Enter interface mode for eth2.
`(config-if)#ip address 2.2.2.1/24`	Configure the IP address for interface eth2 to be in network 0.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode for eth1.
`(config-if)#ip address 1.1.1.1/24`	Configure the IP address for interface eth1 to be in network 1.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.

(config)#router vrrp 1 eth2	Create a VRRP instance for interface eth2.
(config-router)#virtual-ip 2.2.2.1 owner	Configure R1 as the owner.
(config-router)#advertisement-interval 100	Configure the default value for the advertisement interval. The configurable range is 5 to 4095 centi seconds (value must be a multiple of 5).
(config-router)#preempt-mode true	Set `true` as the default value for the field.
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

R2

#configure terminal	Enter the Configure mode.
(config)#interface eth2	Enter interface mode for eth2.
(config-if)#ip address 2.2.2.2/24	Configure the IP address for interface eth2 to be in network 0.
(config-if)#commit	Commit the candidate configuration to the running configuration.
(config-if)#exit	Exit interface mode.
(config)#interface eth1	Enter interface mode for `eth1`.
(config-if)#ip address 1.1.1.2/24	Configure the IP address for interface `eth1` to be in network 1.
(config-if)#commit	Commit the candidate configuration to the running configuration.
(config-if)#exit	Exit interface mode.
(config)#router vrrp 1 eth2	Create a VRRP instance for interface eth2.
(config-router)#virtual-ip 2.2.2.1	Configure R2 as the backup.
(config-router)#advertisement-interval 100	Configure the default value for the advertisement interval. The configurable range is 5 to 4095 centi seconds (value must be a multiple of 5).
(config-router)#priority 100	Set the default value for the backup router.
(config-router)#preempt-mode true	Set `true` as the default value for the field.
(config-router)# v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable the VRRP session on the router.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

R3

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth2`	Enter interface mode for eth2.
`(config-if)#ip address 2.2.2.3/24`	Configure the IP address for interface eth2 to be in network 0.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth1`	Enter interface mode for eth1.
`(config-if)#ip address 1.1.1.3/24`	Configure the IP address for interface eth1 to be in network 1.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode.
`(config)#router vrrp 1 eth2`	Create a VRRP instance for interface eth2.
`(config-router)#virtual-ip 2.2.2.1`	Configure R3 as the backup.
`(config-router)#advertisement-interval 100`	Configure the default value for the advertisement interval. The configurable range is 5 to 4095 centi seconds (value must be a multiple of 5).
`(config-router)#priority 254`	Configure the priority for R3. The configurable range is 1-255.
`(config-router)#preempt-mode true`	Set `true` as the default value for the field.
`(config-router)# v2-compatible`	Enable the v2-compatible
`(config-router)#authentication text abcd`	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
`(config-router)#enable`	Enable the VRRP session on the router.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode.

Validation

Router 1

```
R1#show vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminUp   - Master
 Virtual IP address: 2.2.2.1 (Owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 2.2.2.1
 Operational master IP address: 2.2.2.1
 Priority is 255
 Advertisement interval: 100 centi sec
```

```
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 9 minutes 6 seconds (54600 centi sec)
 Master uptime: 0 hours 1 minutes 18 seconds (7800 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface eth2: JOINED
 V2-Compatible: TRUE
```

Router 2

```
R2#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminUp   - Backup
 Virtual IP address: 2.2.2.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 2.2.2.2
 Operational master IP address: 2.2.2.1
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 11 minutes 28 seconds (68800 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface eth2: JOINED
 V2-Compatible: TRUE
```

Router 3

```
R3#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminUp   - Backup
 Virtual IP address: 2.2.2.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 2.2.2.3
 Operational master IP address: 2.2.2.1
 Priority is 254
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 14 minutes 23 seconds (86300 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
```

```
 Multicast membership on IPv4 interface eth2: JOINED
 V2-Compatible: TRUE
```

Ping Output at Host A

```
HOSTA#ping 1.1.1.5
Press CTRL+C to exit
PING 1.1.1.5 (1.1.1.5) 56(84) bytes of data.
64 bytes from 1.1.1.5: icmp_seq=1 ttl=63 time=1.40 ms
64 bytes from 1.1.1.5: icmp_seq=2 ttl=63 time=1.09 ms
```

Disabling the Master/Owner

`#configure terminal`	Enter the Configure mode.
`(config)#router vrrp 1 eth2`	Create a VRRP instance for interface eth2.
`(config-router)#disable`	Disable the VRRP session.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode.

Output After Disabling the Master

Router 1

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminDown - Init (admin state down)
 Virtual IP address: 2.2.2.1 (Owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 2.2.2.1
 Operational master IP address: 2.2.2.1
 Priority is 255
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 0 minutes 0 seconds (0 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface eth2: LEFT
 V2-Compatible: TRUE
```

Router 3

```
R3#sh vrrp
VRRP Version: 3
```

```
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: eth2
 State: AdminUp   - Master
 Virtual IP address: 2.2.2.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 2.2.2.3
 Operational master IP address: 2.2.2.3
 Priority is 254
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 40 minutes 55 seconds (245500 centi sec)
 Master uptime: 0 hours 1 minutes 18 seconds (7800 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface eth2: JOINED
 V2-Compatible: TRUE

HOSTA#ping 1.1.1.5
Press CTRL+C to exit
PING 1.1.1.5 (1.1.1.5) 56(84) bytes of data.
64 bytes from 1.1.1.5: icmp_seq=1 ttl=63 time=1.40 ms
64 bytes from 1.1.1.5: icmp_seq=2 ttl=63 time=1.09 m
```

Interface Tracking

The need for VRRP Interface Tracking arose because VRRPv3 was unable to track the gateway interface status. The VRRP Interface Tracking feature provides dynamic failover of an entire circuit, in the event that one member of the group fails. It introduces the concept of a circuit, where two or more Virtual Routers on a single system can be grouped. In the event that a failure occurs, and one of the Virtual Routers performs the Master to Backup transition, the other Virtual Routers in the group are notified, and are forced into the Master to Backup transition, so that both incoming and outgoing packets are routed through the same gateway router, eliminating the problem for Firewall/NAT environments.

Note: VRRP Interface Tracking feature allows you to track the state of an upstream Interface and update the VRRP router priority accordingly. Up to 5 interfaces can be tracked per VRRP session for upstream interfaces and Interface tracking will be disabled when the session is the owner of the VIP.

To configure VRRP Interface Tracking, each circuit is configured to have a corresponding priority-delta value, which is passed to VRRP when a failure occurs. The priority of each Virtual Router on the circuit is decremented by the priority-delta value, causing the VR Master to VR Backup transition.

In this example, two routers, R1 and R2, are configured as backup routers with different priorities. The priority-delta value is configured to be greater than the difference of both the priorities. R1 is configured to have a priority of 150, and R2 has a priority of 50. R1, with a greater priority, is the Virtual Router Master. The priority-delta value is 110, greater than 100 (150 minus 50). On R1, when the external interface xe41, xe50/1 and xe50/2 fails, the priority of R1 becomes 40 (150 minus 110). Since R2 has a greater priority (50) than R1, R2 becomes the VR Master, and routing of packages continues without interruption. When this VR Backup (R1) is up again, it regains its original priority (150), and becomes the VR Master again.

Topology

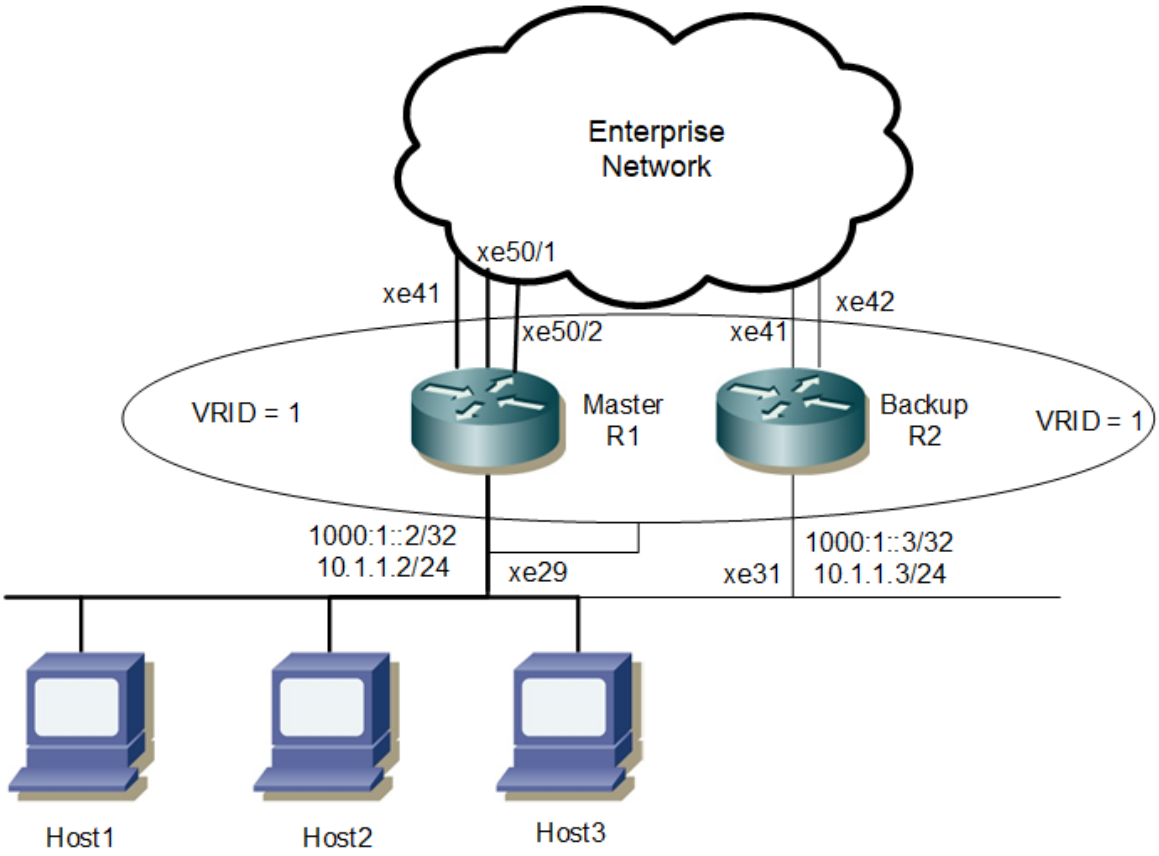

Figure 1-81: VRRP Interface Tracking

R1

`(config)#configure terminal`	Enter the `Configure` mode.
`(config)#bridge 1 protocol rstp vlan-bridge`	Create a RSTP VLAN aware bridge.
`(config)#vlan 1001-1002 bridge 1`	Create VLAN 1001,1002 and map it to bridge 1.
`(config)#interface xe29`	Enter interface mode.
`(config-if)#switchport`	Configure switch port.
`(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group to l2 interface with spanning-tree Disable.
`(config-if)#switchport mode trunk`	Configure switch port mode as trunk.
`(config-if)#switchport trunk allowed vlan add 1001,1002`	Allow vlan 1001 ,1002 on the interface.
`(config-if)#interface vlan1.1001`	Enter the vlan interface
`(config-if)#ip address 10.1.1.2/24`	Configure ipv4 address
`(config-if)#ipv6 address 1000:1::2/32`	Configure ipv6 address
`(config-if)#ipv6 router ospf area 0.0.0.0`	Tag ospfv3 instance on interface with area 0
`(config-if)#exit`	Exit the interface mode
`(config)#router vrrp 1 vlan1.1001`	Create a new ipv4 VRRP instance for interface vlan1.1001

(config-router)#virtual-ip 10.1.1.1	Set the virtual IP address for the VRRP session
(config-router)#priority 150	Configure the priority to 150
(config-router)#circuit-failover xe41 70	Configure the priority-delta value to be 70. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#circuit-failover xe50/1 10	Configure the priority-delta value to be 10. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#circuit-failover xe50/2 30	Configure the priority-delta value to be 10. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters.
(config-router)#enable	Enable the VRRP session on the router.
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode.
(config)#router ipv6 vrrp 1 vlan1.1001	Create a new ipv6 VRRP instance for interface vlan1.1001.
(config-router)#virtual-ipv6 fe80::1	Set the virtual IPv6 address for the VRRP session.
(config-router)#priority 150	Configure the priority to 150
(config-router)#circuit-failover xe41 70	Configure the priority-delta value to be 70. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#circuit-failover xe50/1 10	Configure the priority-delta value to be 10. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#circuit-failover xe50/2 30	Configure the priority-delta value to be 10. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#enable	Enable the VRRP session on the router.
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode
(config)#interface xe50/1	Enter interface mode
(config-if)#ip address 20.1.1.1/24	Configure ipv4 address
(config-if)#ipv6 address 2000::1/64	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0
(config-if)#interface xe50/2	Enter interface mode
(config-if)#ip address 30.1.1.1/24	Configure ipv4 address
(config-if)#ipv6 address 3000::1/64	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0
(config-if)#interface xe41	Enter interface mode
(config-if)#ip address 60.1.1.1/24	Configure ipv4 address
(config-if)#ipv6 address 6000::1/64	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0

(config-if)#commit	Commit the configuration on the node
(config-if)#exit	Exit interface and configure mode

R2

(config)#bridge 1 protocol rstp vlan-bridge	Create a RSTP VLAN aware bridge with bridge-id.
(config)#vlan 1001-1002 bridge 1	Create VLAN 1001,1002 and map it to bridge 1
(config)#interface xe31	Enter interface mode
(config-if)#switchport	Configure switch port
(config-if)#bridge-group 1 spanning-tree disable	Configure bridge group to l2 interface with spanning-tree Disable
(config-if)#switchport mode trunk	Configure switch port mode as trunk
(config-if)#switchport trunk allowed vlan add 1001,1002	Allow vlan 1001 ,1002 on the interface
(config-if)#interface vlan1.1001	Enter the vlan interface
(config-if)#ip address 10.1.1.3/24	Configure ipv4 address
(config-if)#ipv6 address 1000:1::3/32	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0
(config-if)#commit	Commit the configuration on the node
(config-if)#exit	Exit interface mode
(config)#router vrrp 1 vlan1.1001	Create a router ipv4 VRRP instance for interface vlan1.1001.
(config-router)#virtual-ip 10.1.1.1	Set the virtual IP address for the VRRP session.
(config-router)#priority 50	Configure the priority to 50 (less than 150), because R2 is the VR Backup router.
(config-router)#v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters.
(config-router)#enable	Enable the VRRP session
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode
(config)#router ipv6 vrrp 1 vlan1.1001	Create a router ipv6 VRRP ipv6 instance for interface vlan1.1001.
(config-router)#virtual-ipv6 fe80::1	Set the virtual IPv6 address for the VRRP session.
(config-router)#priority 50	Configure the priority to 50 (less than 150), because R2 is the VR Backup router.
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode
(config)#interface xe41	Enter interface mode
(config-if)#ip address 80.1.1.1/24	Configure ipv4 address
(config-if)#ipv6 address 8000::1/64	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0
(config-if)#interface xe42	Enter interface mode
(config-if)#ip address 90.1.1.1/24	Enter the tracked interface

(config-if)#ipv6 address 9000::1/64	Configure ipv6 address
(config-if)#ipv6 router ospf area 0.0.0.0	Tag ospfv3 instance on interface with area 0
(config-if)#commit	Commit the configuration on the node
(config-if)#exit	Exit interface and configure mode

Validation

```
R1# show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv4
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Master
 Virtual IP address: 10.1.1.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.1.1.2
 Operational master IP address: 10.1.1.2
 Configured priority: 150, Current priority: 150
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 1 minutes 12 seconds (7200 centi sec)
 Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Monitored circuit: xe41, Priority Delta: 70, Status: UP
 Monitored circuit: xe50/1, Priority Delta: 10, Status: UP
 Monitored circuit: xe50/2, Priority Delta: 30, Status: UP
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINED V
2-Compatible: TRUE

R1#
R2#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv4
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Backup
 Virtual IP address: 10.1.1.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.1.1.3
 Operational master IP address: 10.1.1.2
 Priority is 50
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 7 hours 52 minutes 53 seconds (2837300 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface vlan1.1001: JOINED
 V2-Compatible: TRUE
```

```
R2#
R1#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv6
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Master
 Virtual IP address: fe80::1 (Not-owner)
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::ba6a:97ff:fe3c:de9d
 Operational master IP address: fe80::ba6a:97ff:fe3c:de9d
 Configured priority: 150, Current priority: 150
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 3 minutes 54 seconds (23400 centi sec)
 Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Monitored circuit: xe41, Priority Delta: 70, Status: UP
 Monitored circuit: xe50/1, Priority Delta: 10, Status: UP
 Monitored circuit: xe50/2, Priority Delta: 30, Status: UP
 Multicast membership on IPv6 interface  vlan1.1001: JOINED
 V2-Compatible: FALSE
R1#
R2#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv6
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Backup
 Virtual IP address: fe80::1 (Not-owner)
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::82a2:35ff:fe35:135f
 Operational master IP address: fe80::ba6a:97ff:fe3c:de9d
 Priority is 50
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 7 hours 55 minutes 11 seconds (2851100 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  vlan1.1001: JOINED
 V2-Compatible: FALSE
R2#
```

After shut down of all tracked interfaces (xe50/1, xe50/2 and xe41) in R1:

```
R1#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv4
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Backup
 Virtual IP address: 10.1.1.1 (Not-owner)
```

```
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.1.1.2
 Operational master IP address: 10.1.1.3
 Configured priority: 150, Current priority: 40
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 7 minutes 46 seconds (46600 centi sec)
 Skew time: 84 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Monitored circuit: xe41, Priority Delta: 70, Status: DOWN
 Monitored circuit: xe50/1, Priority Delta: 10, Status: DOWN
 Monitored circuit: xe50/2, Priority Delta: 30, Status: DOWN
 Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINED
V2-Compatible: TRUE
R1#
R2#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv4
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Master
 Virtual IP address: 10.1.1.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.1.1.3
 Operational master IP address: 10.1.1.3
 Priority is 50
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 7 hours 57 minutes 41 seconds (2866100 centi sec)
 Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv4 interface vlan1.1001: JOINED
 V2-Compatible: FALSE
R2#
R1#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv6
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Backup
 Virtual IP address: fe80::1 (Not-owner)
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::ba6a:97ff:fe3c:de9d
 Operational master IP address: fe80::82a2:35ff:fe35:135f
 Configured priority: 150, Current priority: 40
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 8 minutes 43 seconds (52300 centi sec)
 Skew time: 84 centi sec
 Master Down Interval: 380 centi sec
 Accept mode: TRUE
```

```
 Preempt mode: TRUE
 Monitored circuit: xe41, Priority Delta: 70, Status: DOWN
 Monitored circuit: xe50/1, Priority Delta: 10, Status: DOWN
 Monitored circuit: xe50/2, Priority Delta: 30, Status: DOWN
 Multicast membership on IPv6 interface  vlan1.1001: JOINED
 V2-Compatible: FALSE
R1#
R2#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled
Address family IPv6
VRRP Id: 1 on interface: vlan1.1001
 State: AdminUp   - Master
 Virtual IP address: fe80::1 (Not-owner)
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::82a2:35ff:fe35:135f
 Operational master IP address: fe80::82a2:35ff:fe35:135f
 Priority is 50
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 7 hours 59 minutes 4 seconds (2874400 centi sec)
 Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  vlan1.1001: JOINED
 V2-Compatible: FALSE
R2#
```

VRRP-Backward Compatibility

This section contains VRRP Backward Compatibility configuration examples.

The backward compatibility feature which implements version 3 of VRRP protocol recognizes the presence of VRRP version 2 compatible routers in the network and performs all operations normally. This support is intended for upgrade scenarios and is not recommended for permanent deployments. This should only occur when a router is transitioning from VRRPv2 to VRRPv3.

VRRP Backward Compatibility is applicable only for VRRP IPv4.

Topology

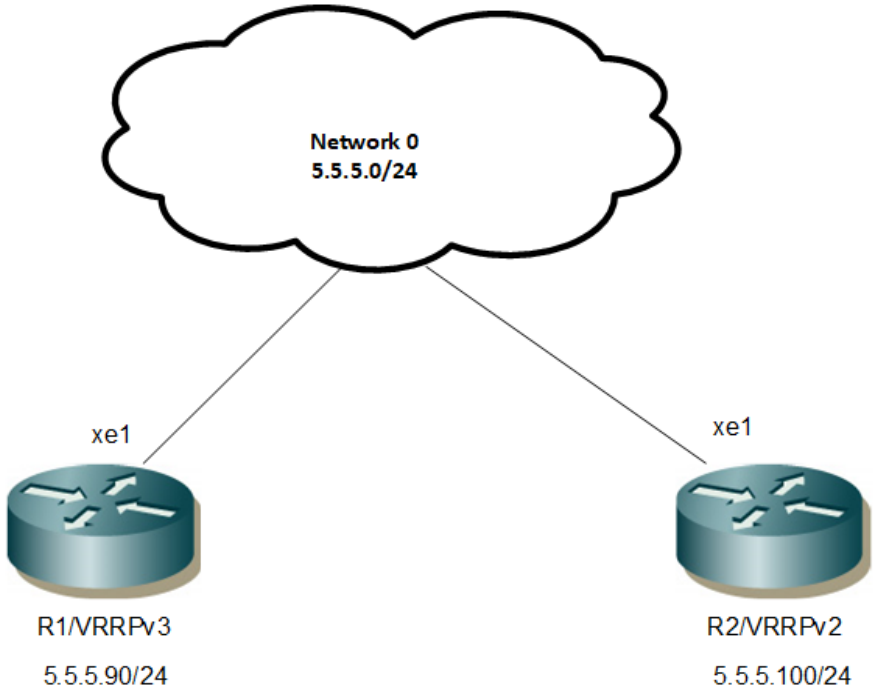

R1

`#configure terminal`	Enter the Configure mode.
`(config)#interface xe1`	Enter Interface configuration mode.
`(config-if)#ip address 5.5.5.90/24`	Configure IP address on the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit Router mode.
`(config)#router vrrp 1 xe1`	Configure VRRP on DUT with Virtual router Identifier as 1 on the interface xe1.
`(config-router)#virtual-ip 5.5.5.190`	Configure Virtual-IP address as the interface IP address of Owner which is not actually present in the LAN.
`(config-router)#v2-compatible`	Enable the v2-compatible
`(config-router)#authentication text abcd`	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
`(config-router)#enable`	Enable VRRP session on DUT.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode.
`(config)#vrrp compatible-v2 enable`	Enable VRRP-Backward compatibility feature on a VRRPv3 running router.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode.

© 2025 IP Infusion Inc. Proprietary

R2

#configure terminal	Enter the Configure mode.
(config)# interface xe1	Enter Interface configuration mode.
(config-if)#ip address 5.5.5.100/24	Configure IP address on the interface.
(config-if)#commit	Commit the candidate configuration to the running configuration.
(config-if)#exit	Exit Router mode.
(config)#router vrrp 1 xe1	Configure VRRP on DUT with Virtual router Identifier as 1 on the interface xe1.
(config-router)#virtual-ip 5.5.5.190	Configure Virtual-IP address as the interface IP address of Owner which is not actually present in the LAN.
(config-router)#v2-compatible	Enable the v2-compatible
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for vrrpv2 packets, accept only 8 characters
(config-router)#enable	Enable VRRP session on DUT.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#exit	Exit Router mode.

Validation

R1

```
OcNOS#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility enabled

Address family IPv4
VRRP Id: 1 on interface: xe1
State: AdminUp  - Master
Virtual IP address: 5.5.5.190 (Not-owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 5.5.5.90
Operational master IP address: 5.5.5.90
Priority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 2 minutes 32 seconds (15200 centi sec)
Master uptime: 0 hours 2 minutes 28 seconds (14800 centi sec)
Accept mode: TRUE
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface xe1: JOINED
V2-Compatible: TRUE
```

R2

```
OcNOS#sh vrrp
```

```
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: xe1
State: AdminUp   - Backup
Virtual IP address: 5.5.5.190 (Not-owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 5.5.5.100
Operational master IP address: 5.5.5.90
Priority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 1 minutes 11 seconds (7100 centi sec)
Skew time: 60 centi sec
Master Down Interval: 360 centi sec
Accept mode: TRUE
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface xe1: JOINED
V2-Compatible: TRUE
```

Redundancy Using VRRP and OSPF: Two Virtual Routers

This example illustrates a configuration of two routers between two end-hosts. R1 and R2 are two virtual routers functioning as backups for each other, with VRRP running on the 10.10.12.0/24 network (LAN), and OSPF running on the 10.10.10.0/24 network (ISP).

Topology

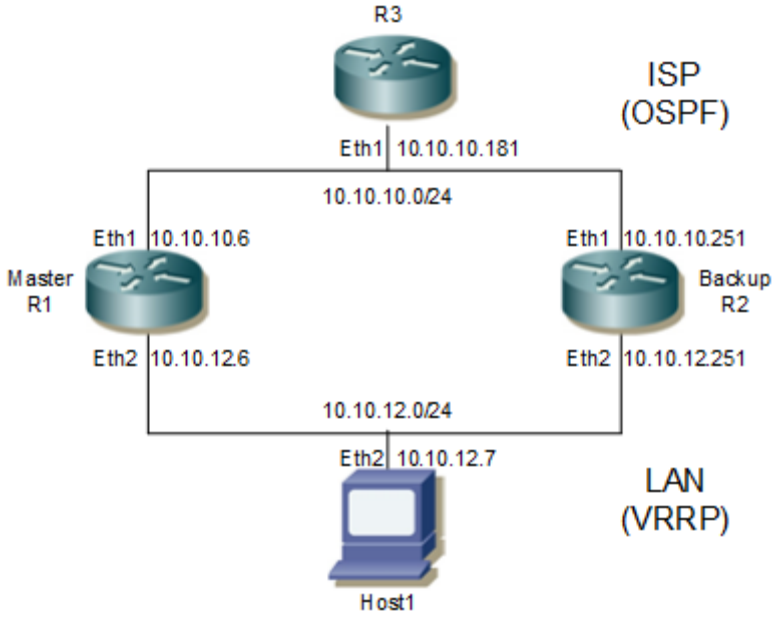

Figure 1-82: Redundancy Using VRRP and OSPF

In Figure 1-82:

- R3 is an OSPF router representing an OSPF network to an ISP.
- R1 is the VRRP Master/OSPF router.
- R2 is the VRRP Backup/OSPF router.
- Host 1 is an end-host.

VRRP handles any failure of the Master's link to the LAN. Failures in the OSPF network that could cause the Master to lose routing information would cause packets from Host 1 that are targeted for R3 to be dropped. Running VRRP on the OSPF network to create redundancy is undesirable, because doing so would cause erroneous VRRP packets to be sent to the ISP.

An alternative method to achieve redundancy is to run OSPF on the LAN side. By running OSPF on the LAN, any routing information lost by the Master would be regained from the Backup on the LAN interface, resulting in ICMP redirects to R2 for traffic received from Host 1. To reduce OSPF control traffic, R1 and R2 are configured as Area Border Routers (ABR), and the LAN is configured as a stub network to reduce LSA advertisement traffic on the LAN. Before enabling OSPF on the LAN, verify that VRRP is running with R1 as the Master and R2 as the Backup.

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: vlan1.1000
State: AdminUp   - Master
Virtual IP address: 10.1.1.2 (Owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.2
Operational master IP address: 10.1.1.2
```

```
Priority is 255
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 1 minutes 11 seconds (7100 centi sec)
Master uptime: 0 hours 1 minutes 11 seconds (7100 centi sec)
Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv4 interface vlan1.1000: JOINED
V2-Compatible: FALSE

R2#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: vlan1.1000
State: AdminUp   - Backup
Virtual IP address: 10.1.1.2 (Not-owner)
Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.3
Operational master IP address: 10.1.1.2
Priority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 0 minutes 38 seconds (3800 centi sec)
Skew time: 60 centi sec
Master Down Interval: 360 centi sec
Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv4 interface vlan1.1000: JOINED
V2-Compatible: FALSE
```

Steps to configure OSPF on the LAN are given below.

R3

`#configure terminal`	Enter the Configure mode.
`(config)#router ospf 1`	Configure the routing process and specify the process ID (`1`). The process ID should be a unique integer.
`(config-router)#ospf router-id 10.10.10.181`	Specify the OSPF router ID.
`(config-router)#timers spf exp 0 0`	Set `timers` to minimum time for quick convergence.
`(config-router)#network 10.10.10.0/24 area 0`	Define one interface (`10.10.10.0/24`) on which OSPF runs and associate the area ID (0).
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit router VRRP mode.

R1

`#configure terminal`	Enter the Configure mode.
`(config)#router ospf 1`	Configure the routing process and specify the process ID (`1`). The process ID should be a unique integer.
`(config-router)#ospf router-id 10.10.12.6`	Specify the OSPF router ID.
`(config-router)#area 1 stub`	Define area 1 as a stub network.
`(config-router)#network 10.10.10.0/24 area 0`	Define one interface (`10.10.10.0/24`) on which OSPF runs and associate the area ID (`0`).
`(config-router)#network 10.10.12.0/24 area 1`	Define the other interface (`10.10.12.0/24`) on which OSPF runs and associate the area ID (`1`)
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit router VRRP mode.

R2

`#configure terminal`	Enter the Configure mode.
`(config)#router ospf 1`	Configure the routing process, and specify the process ID (`1`). The process ID should be a unique integer
`(config-router)#ospf router-id 10.10.12.251`	Specify the OSPF router ID.
`(config-router)#area 1 stub`	Define area 1 as a stub network.
`(config-router)#network 10.10.10.0/24 area 0`	Define one interface (`10.10.10.0/24`) on which OSPF runs and associate the area ID (`0`).
`(config-router)#network 10.10.12.0/24 area 1`	Define the other interface (`10.12.10.0/24`) on which OSPF runs and associate the area ID (`1`)
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit router VRRP mode.

Verification

1. Set gateway on the end-host (statically):

```
(root@host1)#route add -net 10.10.10.0 netmask 255.255.255.0 gw 10.10.12.6
```

2. Verify end-host reachability via traceroute:

```
(root@host1)#traceroute 10.10.10.181
traceroute to 10.10.10.181 (10.10.10.181), 30 hops max, 38 byte packets
1 10.10.12.6 (10.10.12.6) 0.835 ms 0.350 ms 0.341 ms
2 10.10.10.181 (10.10.10.181) 9.557 ms 0.572 ms 0.545 ms
```

3. Bring down eth2 of R1:

```
[root@r1 sbin]#ifconfig eth2 down
```

4. Verify end-host reachability via traceroute:

```
(root@host1)#traceroute 10.10.10.181
traceroute to 10.10.10.181 (10.10.10.181), 30 hops max, 38 byte packets
1 10.10.12.6 (10.10.12.6) 0.461 ms 0.352 ms 0.334 ms
```

```
2 10.10.12.251 (10.10.12.251) 0.425 ms 0.432 ms 0.410 ms
3 10.10.10.181 (10.10.10.181) 0.691 ms 0.639 ms 0.607 ms
```

5. Bring up eth2 of R1:

```
[root@r1 sbin]#ifconfig eth2 up
```

6. Verify end-host reachability via traceroute:

```
(root@host1)#traceroute 10.10.10.181
traceroute to 10.10.10.181 (10.10.10.181), 30 hops max, 38 byte packets
1 10.10.12.6 (10.10.12.6) 0.457 ms 0.356 ms 0.443 ms
2 10.10.10.181 (10.10.10.181) 0.698 ms 0.642 ms 0.618 ms
```

VRRP Over MLAG

This section contains VRRP over MLAG configuration examples.

In this configuration TOR1 and TOR2 forms the VRRP master/backup relationship over MLAG interface.

Topology

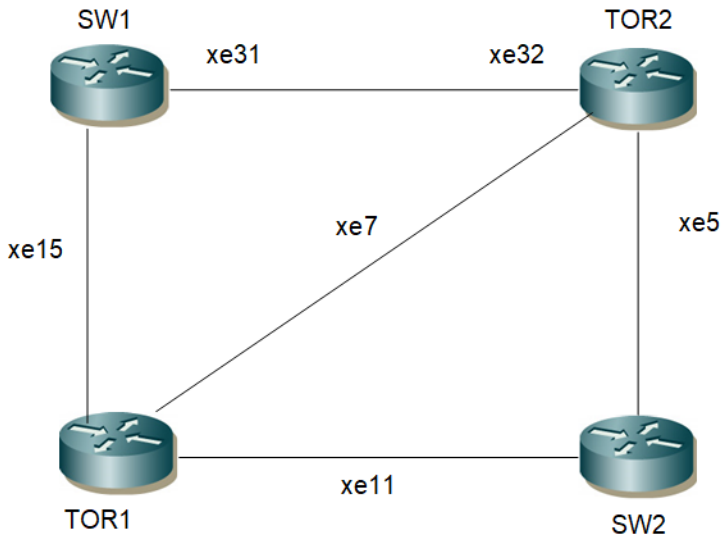

Figure 1-83: VRRP over MLAG

SW1

`#config terminal`	Enter the Configure mode.
`(config)#hostname SW1`	Assign the hostname for the router.
`SW1(config)# bridge 1 protocol rstp vlan-bridge`	Create bridge
`SW1(config)#vlan database`	Enter to VLAN database
`SW1(config-vlan)#vlan 2-1002 bridge 1 state enable`	Create VLANs
`SW1(config-vlan)#exit`	Exit the VLAN database mode
`SW1(config)# int po1`	Enter the interface mode
`SW1(config-if)#switchport`	Configure the interface as Layer 2
`SW1(config-if)# bridge-group 1`	Assign the bridge to the interface.
`SW1(config-if)#switchport mode trunk`	Configure the interface as trunk mode
`SW1(config-if)#switchport trunk allowed vlan all`	Configure the interface to allow all VLAN IDs
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`SW1(config-if)#exit`	Exit the interface mode
`SW1(config)# int xe15`	Enter the interface mode

SW1(config-if)# channel-group 1 mode active	Add the interface as member of LAG interface
SW1(config-if)# int xe31	Enter the interface mode
SW1(config-if)# channel-group 1 mode active	Add the interface as member of LAG interface
SW1(config-if)# int xe36	Enter the interface mode
SW1(config-if)#switchport	Configure the interface as Layer 2
SW1(config-if)# bridge-group 1	Assign the bridge to the interface.
SW1(config-if)#switchport mode trunk	Configure the interface as trunk mode
SW1(config-if)#switchport trunk allowed vlan all	Configure the interface to allow all VLAN IDs
SW1(config-if)#exit	Exit the interface mode
SW1(config-if)#commit	Commit the candidate configuration to the running configuration.
SW1(config-if)#exit	Exit the interface mode.

TOR1

#config terminal	Enter the Config terminal
(config)#hostname TOR1	Assign the hostname to the router
TOR1(config)# bridge 1 protocol rstp vlan-bridge	Configure the bridge
TOR1(config)#vlan database	Enter the VLAN database
TOR1(config-vlan)#vlan 2-1002 bridge 1 state enable	Configure VLANs
TOR1(config-vlan)#commit	Commit the candidate configuration to the running configuration.
TOR1(config-vlan)#exit	Exit the VLAN database
TOR1(config)# int po1	Enter the interface mode
TOR1(config-if)#switchport	Configure the interface as L2
TOR1(config-if)# bridge-group 1	Assign the bridge to the interface
TOR1(config-if)#switchport mode trunk	Configure the interface as trunk mode
TOR1(config-if)#switchport trunk allowed vlan all	Configure to allow all the VLANs
TOR1(config-if)#commit	Commit the candidate configuration to the running configuration.
TOR1(config-if)#exit	Exit the interface mode
TOR1(config)# int xe15	Enter the interface mode
TOR1(config-if)# channel-group 1 mode active	Assign the interface as member of LAG interface
TOR1(config-if)#exit	Exit the interface mode
TOR1(config)#interface xe7	Enter interface mode.
TOR1(config)#switchport	Configure as switchport L2 interface.
TOR1(config)#mcec domain configuration	Enter the MCEC mode
TOR1(config-mcec-domain)#domain-system-number 1	Assign the domain system number for MLAG interface.

`TOR1(config-mcec-domain)#domain-address 1111.2222.3333`	Assign the domain address
`TOR1(config-mcec-domain)#domain-hello-timeout short`	Assign the domain hello timeout
`TOR1(config-mcec-domain)#intra-domain-link xe7`	Assign the interface as IDL
`TOR1(config-mcec-domain)#commit`	Commit the candidate configuration to the running configuration.
`TOR1(config-mcec-domain)#exit`	Exit the MCEC domain
`TOR1(config)#interface lo`	Enter the interface mode
`TOR1(config-if)#ip add 1.1.1.1/32 secondary`	Assign IP address to the interface
`TOR1(config-if)#int xe11`	Enter the interface mode
`TOR1(config-if)#ip add 20.20.20.1/24`	Assign IP address to the interface
`TOR1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`TOR1(config-if)#exit`	Exit the interface mode
`TOR1(config)#router ospf 100`	Configure the routing process and specify the process ID (`100`). The process ID should be a unique integer.
`TOR1(config-router)#network 20.20.20.1/24 area 0`	Add the network to the OSPF process
`TOR1(config-router)#network 1.1.1.1/32 area 0`	Define one interface (`1.1.1.1/32`) on which OSPF runs and associate the area ID (0).
`TOR1(config-router)#redistribute connected`	Redistribute the connected routes
`TOR1(config-router)#commit`	Commit the candidate configuration to the running configuration
`TOR1(config-router)#exit`	Exit the OSPF process
`TOR1(config)#interface mlag1`	Enter mlag1 interface
`TOR1(config)#switchport`	Configure as switchport L2 interface
`TOR1(config)#bridge-group 1`	Assign the bridge to the interface
`TOR1(config)#switchport mode trunk`	Make the interface trunk mode
`TOR1(config)#switchport trunk allowed vlan all`	Configure to allow all vlans
`TOR1(config)#int po1`	Enter the interface mode
`TOR1(config-if)#mlag 1`	Redistribute the connected routes
`TOR1(config-if)#exit`	Exit the interface mode
`TOR1(config)#int vlan1.2`	Enter interface mode
`TOR1(config-if)#ip address 10.10.10.1/24`	Assign IP address
`TOR1(config-if)#exit`	Exit the interface mode
`TOR1(config)#router vrrp 1 vlan1.2`	Create the VRRP process
`TOR1(config-router)# virtual-ip  10.10.10.1 owner`	Assign the virtual IP address to VRRP
`TOR1(config-router)#enable`	Enable the VRRP process
`TOR1(config-router)#commit`	Commit the candidate configuration to the running configuration.
`TOR1(config-router)#exit`	Exit the VRRP mode

TOR2

#config terminal	Enter the Config terminal
(config)#hostname TOR2	Assign the hostname to the router
TOR2(config)# bridge 1 protocol rstp vlan-bridge	Configure the bridge
TOR2(config)#vlan database	Enter the VLAN database
TOR2(config-vlan)#vlan 2-1002 bridge 1 state enable	Configure VLANs
TOR2(config-vlan)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-vlan)#exit	Exit the VLAN database
TOR2(config)# int po1	Enter the interface mode
TOR2(config-if)#switchport	Configure the interface as L2
TOR2(config-if)# bridge-group 1	Assign the bridge to the interface
TOR2(config-if)#switchport mode trunk	Configure the interface as trunk mode
TOR2(config-if)#switchport trunk allowed vlan add 2	Configure to add the VLAN 2 to the receiving packet.
TOR2(config-if)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-if)#exit	Exit the interface mode
TOR2(config)# int xe32	Enter the interface mode
TOR2(config-if)# channel-group 1 mode active	Assign the interface as member of LAG interface
TOR2(config-if)#exit	Exit the interface mode
TOR2(config)#interface xe7	Enter interface mode
TOR2(config)#switchport	Configure as switchport L2 interface
TOR2(config)#mcec domain configuration	Enter the MCEC mode
TOR2(config-mcec-domain)#domain-system-number 2	Assign the domain system number for MLAG interface.
TOR2(config-mcec-domain)#domain-address 1111.2222.3333	Assign the domain address
TOR2(config-mcec-domain)#domain-hello-timeout short	Assign the domain hello timeout
TOR2(config-mcec-domain)#intra-domain-link xe7	Assign the interface as IDL
TOR2(config-mcec-domain)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-mcec-domain)#exit	Exit the MCEC domain
TOR2(config)#interface mlag1	Enter mlag1 interface
TOR2(config)#switchport	Configure as switchport L2 interface
TOR2(config)#bridge-group 1	Assign the bridge to the interface
TOR2(config)#switchport mode trunk	Make the interface trunk mode
TOR2(config)#switchport trunk allowed vlan all	Configure to allow all vlans
TOR2(config)# int po1	Enter the interface mode

TOR2(config-if)# mlag 1	Create the MLAG interface.
TOR2(config-if)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-if)#exit	Exit the interface mode
TOR2(config)#interface lo	Enter the interface mode
TOR2(config-if)#ip add 2.2.2.2/32 secondary	Assign IP address to the interface
TOR2(config-if)#int xe5	Enter the interface mode
TOR2(config-if)#ip add 30.30.30.1/24	Assign IP address to the interface
TOR2(config-if)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-if)#exit	Exit the interface mode
TOR2(config)#router ospf 100	Enter the OSPF router mode
TOR2(config-router)#network 30.30.30.1/24 area 0	Add the network to the OSPF process
TOR2(config-router)#network 2.2.2.2/32 area 0	Add the network to the OSPF process
TOR2(config-router)#redistribute connected	Redistribute the connected routes
TOR2(config-router)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-router)#exit	Exit the OSPF process
TOR2(config)#int vlan1.2	Enter interface mode
TOR2(config-if)#ip address 10.10.10.2/24	Assign IP address
TOR2(config-if)#exit	Exit the interface mode
TOR2(config)#router vrrp 1 vlan1.2	Create the VRRP process
TOR2(config-router)# virtual-ip 10.10.10.1	Assign the virtual IP address to VRRP
TOR2(config-router)#enable	Enable the VRRP process
TOR2(config-router)#commit	Commit the candidate configuration to the running configuration.
TOR2(config-router)#exit	Exit the VRRP mode

SW2

#config terminal	Enter the Config terminal
(config)#hostname SW2	Assign the hostname to the router
SW2(config)# bridge 1 protocol rstp vlan-bridge	Configure the bridge
SW2(config)#vlan database	Enter the VLAN database
SW2(config-vlan)#vlan 2-1002 bridge 1 state enable	Configure VLANs
SW2(config-vlan)#commit	Commit the candidate configuration to the running configuration.
SW2(config-vlan)#exit	Exit the VLAN database
SW2(config)# int xe11	Enter the interface mode
SW2(config-if)#ip address 20.20.20.2/24	Assign IP address to the interface
SW2(config-if)# int xe5	Enter the interface mode

SW2(config-if)#ip address 30.30.30.2/24	Assign IP address to the interface
SW2(config-if)#commit	Commit the candidate configuration to the running configuration.
SW2(config-if)# int xe46	Enter the interface mode
SW2(config-if)#ip address 40.40.40.1/24	Assign IP address to the interface
SW2(config-if)#int lo	Enter the interface mode
SW2(config-if)#ip add 3.3.3.3/32 secondary	Assign IP address to the interface
SW2(config-if)#int xe32	Enter the interface mode
SW2(config-if)#ip add 50.50.50.2/24	Assign IP address to the interface
SW2(config-if)#exit	Exit interface mode
SW2(config)# router ospf 100	Enter OSPF router
SW2(config-router)# network 3.3.3.3/32 area 0	Add the IP address to the OSPF process
SW2(config-router)# network 20.20.20.0/24 area 0	Add the IP address to the OSPF process
SW2(config-router)# network 30.30.30.0/24 area 0	Add the IP address to the OSPF process
SW2(config-router)#network 50.50.50.2/24 area 0	Add the IP address to the OSPF process
SW2(config-router)#commit	Commit the candidate configuration to the running configuration.
SW2(config-router)#exit	Exit the OSPF mode

Validation

TOR1

```
TOR1#show mlag domain summary

-----------------------------------
Domain Configuration
-----------------------------------

Domain System Number            : 1
Domain Address                  : 1111.2222.3333
Domain Priority                 : 32768
Intra Domain Interface          : xe7
Domain Adjacency                : UP
Domain Sync via                 : Intra-domain-interface
-----------------------------------
MLAG Configuration
-----------------------------------

MLAG-1
  Mapped Aggregator             : po1
  Physical properties Digest    : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22
  Total Bandwidth               : 10g
  Mlag Sync                     : IN_SYNC
  Mode                          : Active-Standby
```

```
  Current Mlag state           : Active
  Switchover-mode              : Revertive

TOR1#show vrrp 1 vlan1.2
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: vlan1.2
 State: AdminUp   - Master
 Virtual IP address: 10.10.10.1 (Owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.10.10.1
 Operational master IP address: 10.10.10.1
 Priority is 255
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 5 minutes 11 seconds (31100 centi sec)
 Master uptime: 0 hours 5 minutes 11 seconds (31100 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface vlan1.2: JOINED
 V2-Compatible: TRUE
```

TOR2

```
TOR2#show mlag domain summary

-------------------------------------
Domain Configuration
-------------------------------------

Domain System Number            : 2
Domain Address                  : 1111.2222.3333
Domain Priority                 : 32768
Intra Domain Interface          : xe7
Domain Adjacency                : UP
Domain Sync via                 : Intra-domain-interface
-------------------------------------
MLAG Configuration
-------------------------------------

MLAG-1
  Mapped Aggregator             : po1
  Physical properties Digest    : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22
  Total Bandwidth               : 10g
  Mlag Sync                     : IN_SYNC
  Mode                          : Active-Standby
```

```
  Current Mlag state           : Standby
  Switchover-mode              : Revertive

TOR2#show vrrp 1 vlan1.2
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv4
VRRP Id: 1 on interface: vlan1.2
 State: AdminUp   - Backup
 Virtual IP address: 10.10.10.1 (Not-owner)
 Virtual MAC address is 0000.5e00.0101
 Operational primary IP address: 10.10.10.2
 Operational master IP address: 10.10.10.1
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 6 minutes 27 seconds (38700 centi sec)
 Skew time: 80 centi sec
 Master Down Interval: 380 centi sec
 Preempt mode: TRUE
 Auth-type: simple text, String: abcd
 Multicast membership on IPv4 interface vlan1.2: JOINED
 V2-Compatible: TRUE
 Session is on MLAG interface. Dataplane acting as Master
```

OBJECT Tracking Using IPSLA

This feature is used to track the state of an object for reachability using IP SLA. A client process such as Virtual Router Redundancy Protocol (VRRP), or RIB, registers itself to track objects and then get notified when a state change occurs.

IP SLA (Service-Level Assurance Protocol) is a Performance Measurement protocol. The protocol is used to Analyze IP Service Levels for IP applications and services. IP SLA's uses active traffic-monitoring technology to monitor continuous traffic on the network.

IP SLA uses Internet Control Message Protocol (ICMP) pings to identify a link failure and notifies to the clients that are registered for tracking. IP SLA is supported from OcNOS 5.0.

The Object Tracking feature provides complete separation between the objects to be tracked and the action to be taken by a client when a tracked object state changes. Thus, several clients such as VRRP, or RIB can register their interest with the tracking process, track the same object, and each take different action when the object changes. The Tracking feature is present in OAMD.

Each tracked object is identified by a unique number that is specified on the tracking CLI. Client processes use this number to track a specific object.

The tracking process processes events from the tracked objects and notes any change of value. The changes in the tracked object are communicated to interested client processes, either immediately or after a specified delay. The object values are reported as either up or down.

To configure VRRP Object Tracking, the object is configured to have a priority-delta value, which is passed to VRRP when a failure occurs. The priority of Virtual Router on the circuit is decremented by the priority- delta value, causing the VR Master to VR Backup transition.

In this example, two routers, R1 and R2, are configured as Master and backup routers with different priorities respectively. The priority-delta value is configured to be greater than the difference of both the priorities. R1 is configured to have a priority of 150, and R2 has a priority of 100. R1, with a greater priority, is the Virtual Router Master. The priority-delta value is 60, greater than 50 (150 minus 100). On R1, when the external interface xe41 fails, the Track state becomes DOWN and the priority of R1 becomes 90 (150 minus 60). Since R2 has a greater priority (100) than R1, R2 becomes the VR Master, and routing of packages continues without interruption. When the track state comes UP this VR Backup (R1) is UP again, it regains its original priority (150), and becomes the VR Master again.

Topology

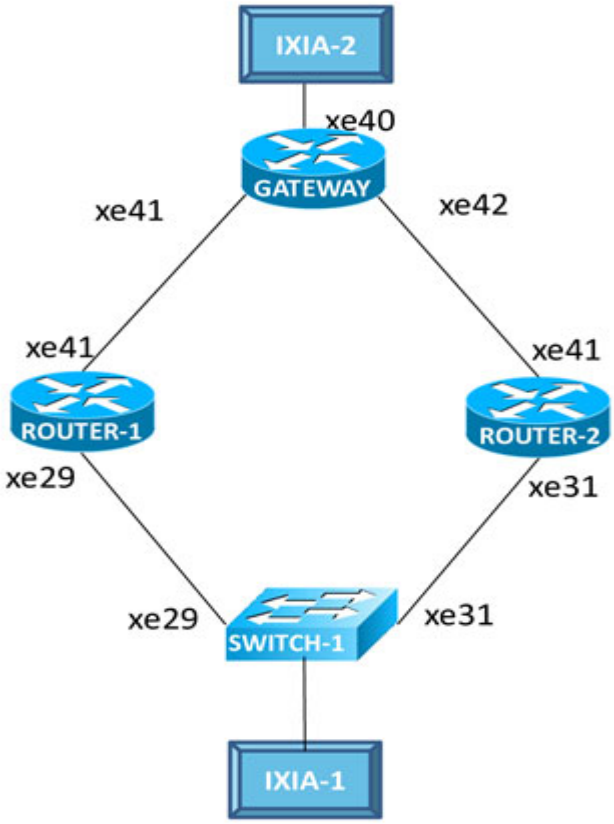

Figure 1-84: VRRP Object Tracking

R1

```
(config)#configure terminal
(config)#bridge 1 protocol rstp vlan-bridge
(config)#vlan 1001-1002 bridge 1
(config)#interface xe29
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree
disable
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add
1001,1002
(config-if)#interface vlan1.1001
(config-if)#ip address 10.1.1.2/24
(config-if)#ipv6 address 1000:1::2/32
(config-if)#ipv6 router ospf area 0.0.0.0
(config-if)#exit
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 4.4.4.4
source-interface xe41
(config)#time-range tr1
(config-tr)#start-time now
(config-tr)#commit
(config-tr)#end-time after 100
(config)#track 1 ip sla 1 reachability
(config)#router vrrp 1 vlan1.1001
(config-router)#virtual-ip 10.1.1.1
(config-router)#priority 150
(config-router)# track 1 decrement 60


(config-router)#v2-compatible
(config-router)#enable
(config-router)#commit
(config-router)#exit
(config)#router ipv6 vrrp 1 vlan1.1001
(config-router)#virtual-ipv6 fe80::1
(config-router)#priority 150
(config-router)# track 1 decrement 60
```

```
(config-router)#enable
(config-router)#commit
(config-router)#exit
(config-if)#interface xe41
(config-if)#ip address 60.1.1.1/24
(config-if)#ipv6 address 6000::1/64
(config-if)#ipv6 router ospf area 0.0.0.0
(config-if)#commit
(config-if)#exit
```

R2

```
(config)#bridge 1 protocol rstp vlan-bridge
(config)#vlan 1001-1002 bridge 1
(config)#interface xe31
(config-if)#switchport
(config-if)#bridge-group 1 spanning-tree
disable
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add
1001,1002
(config-if)#interface vlan1.1001
(config-if)#ip address 10.1.1.3/24
(config-if)#ipv6 address 1000:1::3/32
(config-if)#ipv6 router ospf area 0.0.0.0
(config-if)#commit
(config-if)#exit
(config)#router vrrp 1 vlan1.1001

(config-router)#virtual-ip 10.1.1.1
(config-router)#priority 50

(config-router)#v2-compatible
(config-router)#authentication text abcd

(config-router)#enable
(config-router)#commit
(config-router)#exit
(config)#router ipv6 vrrp 1 vlan1.1001

(config-router)#virtual-ipv6 fe80::1
(config-router)#priority 50

(config-router)#commit
(config-router)#exit
(config)#interface xe41
(config-if)#ip address 80.1.1.1/24
(config-if)#ipv6 address 8000::1/64
(config-if)#ipv6 router ospf area 0.0.0.0
```

Validation

```
R1# sh track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    2 changes, last change : 2019 Feb 19 07:19:57

R1# show vrrp 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: 10.1.1.1 (Not-owner) Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.2 Operational master IP address: 10.1.1.2
Configured priority: 150, Current priority: 150Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 1 minutes 12 seconds (7200 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: UP Monitored circuit: xe50/1,
Priority Delta: 10, Status: UPMonitored circuit: xe50/2, Priority Delta: 30, Status: UP
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINED V2-Compatible: TRUE

R1#
R2#show vrrp 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: 10.1.1.1 (Not-owner)Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.3Operational master IP address: 10.1.1.2Priority
is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 52 minutes 53 seconds (2837300 centi sec)Skew time: 80
centi sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: TRUE



R2#
R1#show vrrp ipv6 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::ba6a:97ff:fe3c:de9dOperational master IP address:
fe80::ba6a:97ff:fe3c:de9dConfigured priority: 150, Current priority: 150 Advertisement
interval: 100 centi sec
```

```
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 3 minutes 54 seconds (23400 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: UP Monitored circuit: xe50/1,
Priority Delta: 10, Status: UPMonitored circuit: xe50/2, Priority Delta: 30, Status: UP
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE
R1#
R2#show vrrp ipv6 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::82a2:35ff:fe35:135fOperational master IP address:
fe80::ba6a:97ff:fe3c:de9dPriority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 55 minutes 11 seconds (2851100 centi sec)Skew time: 80
centi sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE
R2#
```

After shut down the tracked Object (xe41) in R1:

```
R1# sh track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is DOWN
    2 changes, last change : 2019 Feb 19 07:19:57

R1#show vrrp 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: 10.1.1.1 (Not-owner)


Virtual MAC address is 0000.5e00.0101 Operational primary IP address: 10.1.1.2
Operational master IP address: 10.1.1.3
Priority is 90 (Configured Priority is 150)Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 7 minutes 46 seconds (46600 centi sec)Skew time: 84 centi
sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: DOWN Monitored circuit: xe50/1,
Priority Delta: 10, Status: DOWNMonitored circuit: xe50/2, Priority Delta: 30, Status:
DOWNAuth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: TRUE
R1#
```

```
R2#show vrrp 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: 10.1.1.1 (Not-owner)Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.3Operational master IP address: 10.1.1.3Priority
is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 57 minutes 41 seconds (2866100 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: FALSE
R2#
R1#show vrrp ipv6 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::ba6a:97ff:fe3c:de9dOperational master IP
address: fe80::82a2:35ff:fe35:135fPriority is 90 (Configured Priority is 150)
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 8 minutes 43 seconds (52300 centi sec)Skew time: 84 centi
sec
Master Down Interval: 380 centi secAccept mode: TRUE


Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: DOWN Monitored circuit: xe50/1,
Priority Delta: 10, Status: DOWNMonitored circuit: xe50/2, Priority Delta: 30, Status:
DOWNMulticast membership on IPv6 interface vlan1.1001: JOINEDV2-Compatible: FALSE
R1#
R2#show vrrp ipv6 1 vlan1.1001VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::82a2:35ff:fe35:135fOperational master IP address:
fe80::82a2:35ff:fe35:135fPriority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 59 minutes 4 seconds (2874400 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE
R2#
```

Custom VRF Configuration

This section shows how to configure VRRP over MLAG with Custom VRF.

Topology

For topology, refer to VRRP over MLAG.

Configuration

Configure the below configuration on TOR1 and TOR2.

TOR1

`TOR1#config terminal`	Enter the config terminal.
`TOR1(config)#ip vrf vrf1`	Configure VRF1.
`TOR1(config-vrf)#rd 1:1`	Configure rd 1:1.
`TOR1(config-vrf)#route-target both 1:1`	Configure the route target.
`TOR1(config-vrf)#commit`	Commit the candidate configuration to the running configuration.

TOR2

`TOR2#config terminal`	Enter the config terminal.
`TOR2(config)#ip vrf vrf1`	Configure `VRF1`.
`TOR2(config-vrf)#rd 1:1`	Configure `rd 1:1`.
`TOR2(config-vrf)#route-target both 1:1`	Configure the route target.
`TOR2(config-vrf)#commit`	Commit the candidate configuration to the running configuration.
`TOR2(config-vrf)#end`	Returns to privilege mode.

Validation

The below shows the running output for TOR1 and TOR2:

TOR1 and TOR2

```
OcNOS#show running-config
!
! Software version: EC_AS5835-54X-OcNOS-DC-IPBASE-6.4.1.25-Alpha 10/16/2023 08:05:
33
!
!Last configuration change at 05:28:52 UTC Thu Oct 19 2023 by ocnos
!
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
no service password-encryption
!
```

```
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
ip vrf management
!
ip vrf vrf1
 rd 1:1
!
qos enable
!
no ip domain-lookup
ip domain-lookup vrf management
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature ntp vrf management
ntp enable vrf management
feature rsyslog vrf management
!
vlan database
 vlan-reservation 4041-4094
 vlan 2-201 bridge 1 state enable
!
interface mlag1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 2-201
 mtu 9216
 mode active-active
!
interface po1
 switchport
 mlag 1
!
interface po100
 switchport
!
interface ce49
 channel-group 1 mode active
!
interface ce50
!
interface ce51
 description ***Connected to Core***
```

```
 load-interval 30
 ip address 30.30.30.0/31
 ipv6 address 9000::1/64
 mtu 9216
!
interface ce52
!
interface ce53
 channel-group 100 mode active
!
interface ce54
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface vlan1.2
 ip address 2.2.2.2/24
 ipv6 address 2000::2/64
 mtu 9216
!
interface xe1
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
```

```
interface xe11
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
!
interface xe22
!
interface xe23
!
interface xe24
!
interface xe25
!
interface xe26
!
interface xe27
!
interface xe28
!
interface xe29
!
interface xe30
!
interface xe31
!
interface xe32
!
interface xe33
!
interface xe34
!
interface xe35
!
```

```
interface xe36
!
interface xe37
!
interface xe38
!
interface xe39
!
interface xe40
!
interface xe41
!
interface xe42
!
interface xe43
!
interface xe44
!
interface xe45
!
interface xe46
!
interface xe47
!
interface xe48
!
 exit
!
mcec domain configuration
 domain-address 2222.3333.4444
 domain-system-number 2
 intra-domain-link po100
 domain-hello-timeout short
!
router bgp 100
 bgp router-id 200.200.200.200
 timers bgp 3 9
 neighbor 2.2.2.1 remote-as 100
 neighbor 30.30.30.1 remote-as 200
 neighbor 9000::2 remote-as 200
 !
 address-family ipv4 unicast
 redistribute connected
 neighbor 2.2.2.1 activate
 neighbor 2.2.2.1 next-hop-self
 neighbor 30.30.30.1 activate
 exit-address-family
 !
 address-family ipv6 unicast
 redistribute connected
```

```
 neighbor 9000::2 activate
 exit-address-family
 !
 exit
!
router vrrp 1 vlan1.2
 virtual-ip 2.2.2.1
 enable
!
router ipv6 vrrp 1 vlan1.2
 virtual-ipv6 fe80::1
 virtual-ipv6 2000::3
 enable
!
!
end
```

CHAPTER 2 Object Tracking Using IP SLA

Overview

Object Tracking using IP SLA feature tracks the state of an object for reachability using IP SLA. A client process, such as Virtual Router Redundancy Protocol (VRRP) or RIB, registers itself to track objects and receives notifications when a state change occurs.

IP SLA (Service-Level Assurance Protocol) is a Performance Measurement protocol used to analyze IP service levels for IP applications and services. IP SLA employs active traffic-monitoring technology to monitor network traffic continuously.

IP SLA utilizes Internet Control Message Protocol (ICMP) pings to identify link failures and notify registered clients responsible for tracking.

The Object Tracking feature offers complete separation between tracked objects and the actions taken by clients when a tracked object's state changes. Multiple clients, such as VRRP or RIB, can register their interest in the tracking process, monitor the same object, and take different actions when the object changes state. The Tracking feature is present in OAMD (Object-Action Mapping Database).

Each tracked object is identified by a unique number specified in the tracking CLI. Client processes use this number to monitor a specific object.

The tracking process monitors events from the tracked objects, notes any changes in value, and communicates these changes to interested client processes immediately or after a specified delay. The object values are reported as either `up` or `down`.

To configure VRRP Object Tracking, the object is set to have a priority-delta value, which is passed to VRRP when a failure occurs. This decrement in the priority of the Virtual Router on the circuit triggers a transition from VR Master to VR Backup. The VRRP Object tracking is VRF independent.

Note: If both BFD and track are enabled for a static route, the route won't be installed if either BFD or the tracked object is down.

Topology

In this example, two routers, ROUTER-1 (R1) and ROUTER-2 (R2), are configured as Master and backup routers, each with different priorities. The priority-delta value is set to be greater than the difference between both priorities.

R1 is configured with a priority of 150, while R2 is assigned a priority of 100. Since R1 has the higher priority, it initially serves as the Virtual Router Master. The priority-delta value is set to 60, which is greater than the difference between their priorities (150 minus 100).

When the external interface xe41 on R1 fails, the Track state changes to `DOWN`, and R1's priority decreases to 90 (150 minus 60). However, as R2 still maintains a higher priority (100), it assumes the role of VR Master, ensuring uninterrupted packet routing.

Subsequently, when the track state returns to `UP`, the VR Backup (R1) regains its original priority of 150, reclaiming the position of VR Master. This transition allows for seamless failover and network continuity.

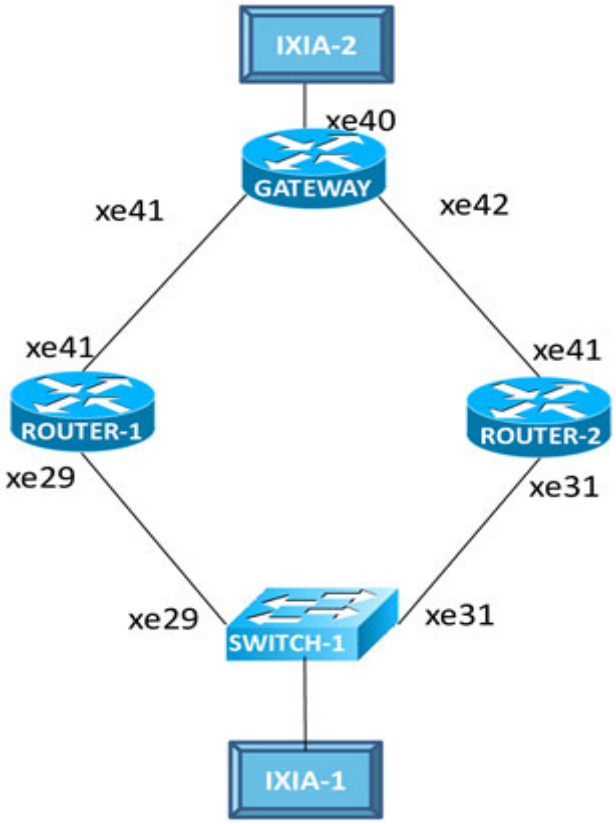

Figure 2-85: VRRP Object Tracking

R1

`(config)#configure terminal`	Enter configure mode.
`(config)#bridge 1 protocol rstp vlan-bridge`	Create a RSTP VLAN aware bridge with bridge-id.
`(config)#vlan 1001-1002 bridge 1`	Create VLAN 1001,1002 and map it to bridge 1.
`(config)#interface xe29`	Enter interface mode.
`(config-if)#switchport`	Configure switch port.
`(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group to L2 interface with spanning-tree disabled.
`(config-if)#switchport mode trunk`	Configure switchport mode as trunk.
`(config-if)#switchport trunk allowed vlan add 1001,1002`	Allow VLAN 1001, and 1002 on the interface.
`(config-if)#interface vlan1.1001`	Enter the VLAN interface.
`(config-if)#ip address 10.1.1.2/24`	Configure IPv4 address.
`(config-if)#ipv6 address 1000:1::2/32`	Configure IPv6 address.
`(config-if)#ipv6 router ospf area 0.0.0.0`	Tag OSPFv3 instance on the interface with area 0.
`(config-if)#exit`	Exit interface mode.
`(config)#ip sla 1`	Configure IP SLA .

(config-ip-sla)#icmp-echo ipv4 4.4.4.4 source-interface xe41	Configure ICMP-echo with a destination IPv4/IPv6 address.
(config-ip-sla-echo)#exit	Exit IP SLA echo mode.
(config)#time-range tr1	Configure a time-range.
(config-tr)#start-time now	Configure a start-time.
(config-tr)#commit	Commit the candidate configuration to the running configuration.
(config-tr)#end-time after 200	Configure end-time.
(config-tr)#exit	Exit time-range mode.
(config)#track 1 ip sla 1 reachability	Configure the track ID for IP SLA.
(config-object-track)#exit	Exit object track mode.
(config)#router vrrp 1 vlan1.1001	Create a router IPv4 VRRP instance for interface vlan1.1001.
(config-router)#virtual-ip 10.1.1.1	Set the virtual IP address for the VRRP session.
(config-router)#priority 150	Configure the priority to 150.
(config-router)#track 1 decrement 60	Configure the priority-delta value to be 60. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#v2-compatible	Enable the v2-compatible.
(config-router)#enable	Enable the VRRP session.
(config-router)#commit	Commit the configuration on the node.
(config-router)#exit	Exit the router mode.
(config)#router ipv6 vrrp 1 vlan1.1001	Create a router IPv6 VRRP IPv6 instance for interface vlan1.1001.
(config-router)#virtual-ipv6 fe80::1	Set the virtual IPv6 address for the VRRP session.
(config-router)#priority 150	Configure the priority to 150.
(config-router)#track 1 decrement 60	Configure the priority-delta value to be 60. In case of failover, this priority-delta value is subtracted from the current VR Master.
(config-router)#enable	Enable the VRRP session.
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode
(config)#interface xe41	Enter interface mode
(config-if)#ip address 60.1.1.1/24	Configure IPv4 address.
(config-if)#ipv6 address 6000::1/64	Configure IPv6 address.
(config-if)#ipv6 router ospf area 0.0.0.0	Tag OSPFv3 instance on the interface with area 0.
(config-if)#commit	Commit the configuration on the node.
(config-if)#exit	Exit interface mode.
(config)#exit	Exit configure mode.

R2

(config)#configure terminal	Enter configure mode.
(config)#bridge 1 protocol rstp vlan-bridge	Create a RSTP VLAN aware bridge with bridge-id.
(config)#vlan 1001-1002 bridge 1	Create VLAN 1001,1002 and map it to bridge 1.
(config)#interface xe31	Enter interface mode.
(config-if)#switchport	Configure switchport.
(config-if)#bridge-group 1 spanning-tree disable	Configure bridge group to L2 interface with spanning-tree disabled.
(config-if)#switchport mode trunk	Configure switchport mode as trunk.
(config-if)#switchport trunk allowed vlan add 1001,1002	Allow VLAN 1001,1002 on the interface.
(config-if)#interface vlan1.1001	Enter the VLAN interface.
(config-if)#ip address 10.1.1.3/24	Configure IPv4 address.
(config-if)#ipv6 address 1000:1::3/32	Configure IPv6 address.
(config-if)#ipv6 router ospf area 0.0.0.0	Tag OSPFv3 instance on interface with area 0.
(config-if)#commit	Commit the configuration on the node.
(config-if)#exit	Exit interface mode.
(config)#router vrrp 1 vlan1.1001	Create a router IPv4 VRRP instance for interface vlan1.1001.
(config-router)#virtual-ip 10.1.1.1	Set the virtual IP address for the VRRP session.
(config-router)#priority 50	Configure the priority to 50 (less than 150), because R2 is the VR Backup router.
(config-router)#v2-compatible	Enable the v2-compatible.
(config-router)#authentication text abcd	Configure the authentication text to specify that as simple text for VRRPv2 packets, accept only 8 characters.
(config-router)#enable	Enable the VRRP session.
(config-router)#commit	Commit the configuration on the node.
(config-router)#exit	Exit the router mode.
(config)#router ipv6 vrrp 1 vlan1.1001	Create a router IPv6 VRRP IPv6 instance for interface vlan1.1001.
(config-router)#virtual-ipv6 fe80::1	Set the virtual IPv6 address for the VRRP session.
(config-router)#priority 50	Configure the priority to 50 (less than 150), because R2 is the VR Backup router.
(config-router)#commit	Commit the configuration on the node
(config-router)#exit	Exit the router mode
(config)#interface xe41	Enter interface mode
(config-if)#ip address 80.1.1.1/24	Configure IPv4 address.
(config-if)#ipv6 address 8000::1/64	Configure IPv6 address.
(config-if)#ipv6 router ospf area 0.0.0.0	Tag OSPFv3 instance on interface with area 0.
(config-if)#exit	Exit interface mode.
(config)#exit	Exit configure mode.

Validation

```
R1#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    2 changes, last change : 2019 Feb 19 07:19:57

R1#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: 10.1.1.1 (Not-owner) Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.2 Operational master IP address: 10.1.1.2
Configured priority: 150, Current priority: 150Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 1 minutes 12 seconds (7200 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec) Accept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: UP Monitored circuit: xe50/1,
Priority Delta: 10, Status: UPMonitored circuit: xe50/2, Priority Delta: 30, Status: UP
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINED V2-Compatible: TRUE

R2#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: 10.1.1.1 (Not-owner)Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.3Operational master IP address: 10.1.1.2Priority
is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 52 minutes 53 seconds (2837300 centi sec)Skew time: 80
centi sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Auth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: TRUE

R1#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::ba6a:97ff:fe3c:de9dOperational master IP address:
fe80::ba6a:97ff:fe3c:de9dConfigured priority: 150, Current priority: 150 Advertisement
interval: 100 centi sec
```

```
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 3 minutes 54 seconds (23400 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec) Accept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: UP Monitored circuit: xe50/1,
Priority Delta: 10, Status: UPMonitored circuit: xe50/2, Priority Delta: 30, Status: UP
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE

R2#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::82a2:35ff:fe35:135fOperational master IP address:
fe80::ba6a:97ff:fe3c:de9dPriority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 55 minutes 11 seconds (2851100 centi sec)Skew time: 80
centi sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE
```

After shut down the tracked Object (xe41) in R1:

```
R1#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is DOWN
    2 changes, last change : 2019 Feb 19 07:19:57

R1#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: 10.1.1.1 (Not-owner)


Virtual MAC address is 0000.5e00.0101 Operational primary IP address: 10.1.1.2
Operational master IP address: 10.1.1.3
Priority is 90 (Configured Priority is 150)Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 7 minutes 46 seconds (46600 centi sec)Skew time: 84 centi
sec
Master Down Interval: 380 centi secAccept mode: TRUE
Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: DOWN Monitored circuit: xe50/1,
Priority Delta: 10, Status: DOWNMonitored circuit: xe50/2, Priority Delta: 30, Status:
DOWNAuth-type: simple text, String: abcd
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: TRUE
```

```
R2#show vrrp 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv4
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: 10.1.1.1 (Not-owner)Virtual MAC address is 0000.5e00.0101
Operational primary IP address: 10.1.1.3Operational master IP address: 10.1.1.3Priority
is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 57 minutes 41 seconds (2866100 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv4 interface vlan1.1001: JOINEDV2-Compatible: FALSE

R1#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Backup
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::ba6a:97ff:fe3c:de9dOperational master IP
address: fe80::82a2:35ff:fe35:135fPriority is 90 (Configured Priority is 150)
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 0 hours 8 minutes 43 seconds (52300 centi sec)Skew time: 84 centi
sec
Master Down Interval: 380 centi secAccept mode: TRUE



Preempt mode: TRUE
Monitored circuit: xe41, Priority Delta: 70, Status: DOWN Monitored circuit: xe50/1,
Priority Delta: 10, Status: DOWNMonitored circuit: xe50/2, Priority Delta: 30, Status:
DOWNMulticast membership on IPv6 interface vlan1.1001: JOINEDV2-Compatible: FALSE

R2#show vrrp ipv6 1 vlan1.1001
VRRP Version: 3
VMAC enabled
Backward Compatibility disabledAddress family IPv6
VRRP Id: 1 on interface: vlan1.1001State: AdminUp- Master
Virtual IP address: fe80::1 (Not-owner)Virtual MAC address is 0000.5e00.0201
Operational primary IP address: fe80::82a2:35ff:fe35:135fOperational master IP address:
fe80::82a2:35ff:fe35:135fPriority is 100
Advertisement interval: 100 centi sec
Master Advertisement interval: 100 centi sec
Virtual router uptime: 7 hours 59 minutes 4 seconds (2874400 centi sec)
Master uptime: 0 hours 2 minutes 34 seconds (15400 centi sec)Accept mode: TRUE
Preempt mode: TRUE
Multicast membership on IPv6 interface  vlan1.1001: JOINEDV2-Compatible: FALSE
```

CHAPTER 3 VRRP IPv6 Configuration

This chapter contains a Virtual Router Redundancy Protocol IPv6 (VRRPv6) configuration example.

VRRPv6 eliminates the risk of a single point of failure inherent in a static default routing environment. It specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRPv6 routers on a LAN. In this sample, OSPFv3 is enabled on Router 1 (R1), the master router, and the backup router. In R1, the connected routes are redistributed.

Note: IPv6 VRRP can be configured using link local address along with additional global ipv6 address as Virtual IP.

Topology

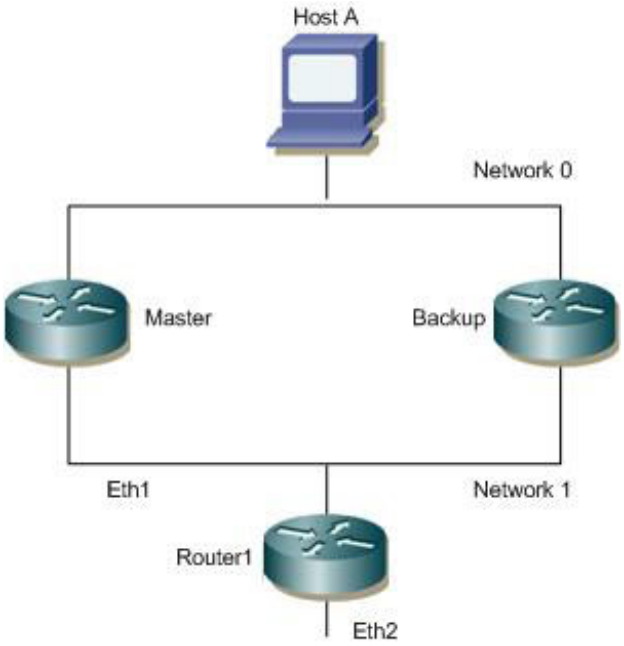

Figure 3-86: Topology

Configuration

Owner/Master Router

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.
`(config-if)#ipv6 address fe80::3037:3aff:fe3a:3a32/64`	Configure the IPv6 address for interface eth0 to be in network 0.
`(config-if)#exit`	Exit the Interface mode.
`(config)#interface eth1`	Enter the Interface mode for eth1.
`(config-if)#ipv6 address fe80::3037:a0ff:fea4:3539/64`	Configure the IPv6 address for interface eth1 to be in network 1.
`(config-if)#exit`	Exit the Interface mode.

`(config)#router ipv6 vrrp 1 eth0`	Create a VRRPv6 instance for interface eth0.
`(config-router)#virtual-ipv6 fe80::3037:3aff:fe3a:3a32 owner`	Configure R1 as the owner.
`(config-router)#advertisement-interval 100`	Configure the default value for the advertisement interval. The configurable range is <5-4095> (value must be a multiple of 5).
`(config-router)#preempt-mode true`	Set true as the default value for the field.
`(config-router)#enable`	Enable the VRRPv6 session on the router.
`(config-router)#exit`	Exit Router mode.

Backup Router

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.
`(config-if)#ipv6 address fe80::3037:3aff:fe3a:3b45/64`	Configure the IPv6 address for interface eth0 to be in network 0.
`(config-if)#exit`	Exit the Interface mode.
`(config)#interface eth1`	Enter the Interface mode for eth1.
`(config-if)#ipv6 address fe80::3037:a0ff:fea4:3a40/64`	Configure the IPv6 address for interface eth1 to be in network 1.
`(config-if)#exit`	Exit the Interface mode.
`(config)#router ipv6 vrrp 1 eth0`	Create a VRRPv6 instance for interface eth0.
`(config-router)#virtual-ipv6 fe80::3037:3aff:fe3a:3a32`	Configure Router 2 (R2) as the backup.
`(config-router)#advertisement-interval 100`	Configure the default value for the advertisement interval. The configurable range is <5-4095> (value must be a multiple of 5).
`(config-router)#priority 100`	Set the default value for the backup router.
`(config-router)#preempt-mode true`	Set `true` as the default value for the field.
`(config-router)#enable`	Enable the VRRPv6 session on the router.
`(config-router)#exit`	Exit Router mode.

Host A

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.
`(config-if)#ipv6 address fe80::3037:3aff:fe3a:3a11/64`	Configure the IPv6 address for interface eth0 to be in network 0.
`(config-if)#exit`	Exit the Interface mode.
`(config)#ipv6 route 5ffe:14:14:14::/64 fe80::3037:3aff:fe3a:3a32`	Configure a static route to reach interface eth2 of R1 through the virtual IPv6 address.

Router 1

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth1`	Enter the Interface mode for eth1.
`(config-if)#ipv6 address fe80::3037:a0ff:fea4:1111/64`	Configure the IPv6 address for interface eth1 to be in network 1.
`(config-if)#exit`	Exit the Interface mode.
`(config)#interface eth2`	Enter the Interface mode for eth2.
`(config-if)#ipv6 address fe80::3437:30ff:fe35:a6ac/64`	Configure the IPv6 address of interface eth2.

Validation

Master Router

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Master
 Virtual IP address: fe80::3037:3aff:fe3a:3a32 (Owner)
 Priority is 255
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Skew time: 0 centi sec
 Accept mode: FALSE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  eth0: JOINED
```

Backup Router

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Backup
 Virtual IP address: fe80::3037:3aff:fe3a:3a32 (Not-Owner)
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Skew time: 0 centi sec
 Accept mode: FALSE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  eth0: JOINED
```

Ping Output at Host A

```
[[root@HstA root]#ping ipv6 fe80::3437:30ff:fe35:a6ac
PING fe80::3437:30ff:fe35:a6ac(fe80::3437:30ff:fe35:a6ac) 56 data bytes
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=1 ttl=63 time=0.398 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=2 ttl=63 time=0.230 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=3 ttl=63 time=0.234 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=4 ttl=63 time=0.230 ms
```

Disabling the Master

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.
`(config-router)#shutdown`	Shut down the interface

Validation

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Init (interface is not running)
 Virtual IP address: fe80::3037:3aff:fe3a:3a32 (Owner)
 Priority is 255
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Skew time: 0 centi sec
 Accept mode: FALSE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface eth0: LEFT
```

Backup Router

```
R1#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Backup
 Virtual IP address: fe80::3037:3aff:fe3a:3a32 (Not-Owner)
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Skew time: 0 centi sec
```

```
 Accept mode: FALSE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  eth0: LEFT
```

Ping Output at Host A

```
[root@HstA root]#ping ipv6 fe80::3437:30ff:fe35:a6ac
PING fe80::3437:30ff:fe35:a6ac(fe80::3437:30ff:fe35:a6ac) 56 data bytes
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=1 ttl=63 time=0.423 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=2 ttl=63 time=0.291 ms
```

GLOBAL IPV6 SUPPORT FOR VRRP

Master Router

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.
`(config-if)#ipv6 address 2001::1/64`	Configure global IPv6 address for interface eth0 to be in network 0.
`(config-if)#commit`	Commit the configurations.
`(config-if)#exit`	Exit the Interface mode.
`(config)#interface eth1`	Enter the Interface mode for eth1.
`(config-if)#ipv6 address 3000::1/64`	Configure global IPv6 address for interface eth1 to be in network 1.
`(config-if)#commit`	Commit the configurations.
`(config-if)#exit`	Exit the Interface mode.
`(config)#router ipv6 vrrp 1 eth0`	Create a VRRPv6 instance for interface eth0.
`(config-router)#virtual-ipv6 fe80::1`	Configure link local address as primary virtual-ip.
`(config-router)#virtual-ipv6 2001::3`	Configure global ipv6 address as additional virtual-ip.
`(config-router)#priority 250`	Set `priority for master node as 250.`
`(config-router)#enable`	Enable the VRRPv6 session on the router.
`(config-router)#commit`	Commit the configurations.
`(config-router)#exit`	Exit Router mode.
`(config)#ipv6 route 4000::/64 3000::2`	Configure static route to reach eth2 of router1.
`(config-router)#commit`	Commit the configurations.
`(config-router)#exit`	Exit Router mode.

Backup Router

`#configure terminal`	Enter the Configure mode.
`(config)#interface eth0`	Enter the Interface mode for eth0.

(config-if)#ipv6 address 2001::2/64	Configure global IPv6 address for interface eth0 to be in network 0.
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.
(config)#interface eth1	Enter the Interface mode for eth1.
(config-if)#ipv6 address 3001::1/64	Configure global IPv6 address for interface eth1 to be in network 1.
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.
(config)#router ipv6 vrrp 1 eth0	Create a VRRPv6 instance for interface eth0.
(config-router)#virtual-ipv6 fe80::1	Configure link local address as primary virtual-ip.
(config-router)#virtual-ipv6 2001::3	Configure global ipv6 address as additional virtual-ip.
(config-router)#priority 100	Set the default value for the backup router.
(config-router)#enable	Enable the VRRPv6 session on the router.
(config-router)#commit	Commit the configurations.
(config-router)#exit	Exit Router mode.
(config)#ipv6 route 4000::/64 3001::2	Configure static route to reach eth2 of router1.
(config-router)#exit	Exit Router mode.

Host A

#configure terminal	Enter the Configure mode.
(config)#bridge 1 protocol rstp vlan-bridge	Create a RSTP VLAN aware bridge with bridge-id.
(config)#vlan 100 bridge 1	Create VLAN 100 and map it to bridge 1
(config)#interface eth0	Enter the Interface mode for eth0.
(config-if)#switchport	Configure switch port.
(config-if)#bridge-group 1 spanning-tree disable	Configure bridge group to l2 interface with spanning-tree Disable
(config-if)#switchport mode trunk	Configure switch port mode as trunk
(config-if)#switchport trunk allowed vlan add 100	Allow vlan 100 on the interface
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode
(config)#interface eth1	Enter the Interface mode for eth1.
(config-if)#switchport	Configure switch port.
(config-if)#bridge-group 1 spanning-tree disable	Configure bridge group to l2 interface with spanning-tree Disable
(config-if)#switchport mode trunk	Configure switch port mode as trunk

(config-if)#switchport trunk allowed vlan add 100	Allow vlan 100 on the interface
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.

Router 1

#configure terminal	Enter the Configure mode.
(config)#interface eth0	Enter the Interface mode for eth0.
(config-if)#ipv6 address 3000::2/64	Configure global IPv6 address for interface eth1 to be in network 1.
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.
(config)#interface eth1	Enter the Interface mode for eth1.
(config-if)#ipv6 address 3001::2/64	Configure global IPv6 address for interface eth1 to be in network 1.
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.
(config)#interface eth2	Enter the Interface mode for eth2.
(config-if)#ipv6 address 4000::1/64	Configure global ipv6 address for interface eth2.
(config-if)#commit	Commit the configurations.
(config-if)#exit	Exit the Interface mode.
(config)#ipv6 route 2001::/64 3001::1	Configure static route to reach eth2 of master.
(config)#ipv6 route 2001::/64 3001::1	Configure static route to reach eth2 of backup.
(config-if)#commit	Commit the configurations.
(config-if)# exit	Exit the Interface mode.

Validation

Master Router

```
OcNOS#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Master
```

```
 Virtual IP address: fe80::1 (Not-owner)
 Virtual IP address: 2001::3
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::923c:b3ff:fe82:8d88
 Operational master IP address: fe80::923c:b3ff:fe82:8d88
 Priority is 250
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 1 minutes 0 seconds (6000 centi sec)
 Master uptime: 0 hours 0 minutes 57 seconds (5700 centi sec)
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  eth0: JOINED
 V2-Compatible: FALSE
```

Backup Router

```
OcNOS#sh vrrp
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Address family IPv6
VRRP Id: 1 on interface: eth0
 State: AdminUp   - Backup
 Virtual IP address: fe80::1 (Not-owner)
 Virtual IP address: 2001::3
 Virtual MAC address is 0000.5e00.0201
 Operational primary IP address: fe80::1644:8fff:fe8e:32e7
 Operational master IP address: fe80::923c:b3ff:fe82:8d88
 Priority is 100
 Advertisement interval: 100 centi sec
 Master Advertisement interval: 100 centi sec
 Virtual router uptime: 0 hours 18 minutes 57 seconds (113700 centi sec)
 Skew time: 60 centi sec
 Master Down Interval: 360 centi sec
 Accept mode: TRUE
 Preempt mode: TRUE
 Multicast membership on IPv6 interface  eth0: JOINED
 V2-Compatible: FALSE
```

Ping From master to backup

```
[[root@HstA root]#ping ipv6 2001::2
PING 2001::2(2001::2) 56 data bytes
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=1 ttl=63 time=0.398 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=2 ttl=63 time=0.230 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=3 ttl=63 time=0.234 ms
64 bytes from fe80::3437:30ff:fe35:a6ac: icmp_seq=4 ttl=63 time=0.230 ms
```

Virtual Router Redundancy Protocol Command Reference

© 2025 IP Infusion Inc. Proprietary

CHAPTER 1 VRRP Commands

This chapter describes the commands for VRRP.

- accept-mode
- advertisement-interval
- authentication text
- circuit-failover
- debug vrrp
- disable
- enable
- operational-ip
- preempt-mode
- priority
- router vrrp
- show debugging vrrp
- show running-config router vrrp
- show vrrp
- show vrrp summary
- show vrrp <1-255>
- show vrrp (global | ipv4) statistics
- snmp restart vrrp
- switch-back-delay
- virtual-ip
- vrrp compatible-v2
- vrrp vmac

accept-mode

Use this command to enable accept mode for the session

Use the `no` parameter with this command to restore the default setting.

Command Syntax

```
accept-mode true
accept-mode false
```

Parameter

None

Default

By default, accept mode for the session is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to set and unset the accept mode.

```
#configure terminal
(config)#router vrrp 2 eth0
(config-router)#accept-mode false
```

advertisement-interval

Use this command to configure the advertisement interval of a virtual router. This is the length of time, in seconds, between each advertisement sent from the master to its backup(s). The master virtual router sends VRRP advertisements to other VRRP routers in the same group. The advertisements communicate the priority and state of the master virtual router. The VRRP advertisements are encapsulated in IP packets and sent to the multicast address assigned to the VRRP group (224.0.0.18). Advertisements are sent every second by default.

Note: VRRP Master router and backup routers should be configured with the same advertisement interval. If there is a mismatch in the configuration, VRRP goes to the INIT state.

Use the `no` parameter with this command to restore the default setting.

Command Syntax

```
advertisement-interval <5-4095>
no advertisement-interval
```

Parameter

`<5-4095>`	Specify the advertisement interval in centi-seconds (in multiples of 5) when VRRPv3 is enabled

Default

By default, advertisements are sent every second

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to configure an advertisement interval of 50 centi-seconds for the virtual router with VR ID 2 on interface eth0.

```
#configure terminal
(config)#router ip vrrp 2 eth0
(config-router)#advertisement-interval 50
```

authentication text

Authenticates VRRPv2 packets received from other routers in the group. If configured authentication, all routers within the VRRP group must use the same authentication string. Authentication is supported only for VRRPv2 packets. V2-compatibility must be enabled.

If both master and backup having same authentication, then one node act as master and another node act as a backup based on priority. If both master and backup having different authentication, both nodes acting as a master.

Use the `no` parameter with this command to remove an authentication.

Command Syntax

```
authentication text TEXT-STRING
no authentication text
```

Parameter

`TEXT-STRING`	Password, maximum eight alphanumeric characters

Default

By default, authentication is disabled.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#configure terminal
(config)#router vrrp 1 eth3
(config-router)#v2-compatible
(config-router)#authentication text abc_123
(config-router)#end

#configure terminal
(config)#router vrrp 1 eth3
(config-router)#no authentication text
(config-router)#end
```

circuit-failover

Use this command to enable the VRRP circuit-failover feature.

Use the `no` form of this command to disable this feature.

Command Syntax

```
circuit-failover IFNAME <1-253>
no circuit-failover (IFNAME|)
```

Parameters

`IFNAME`	Interface of the router that is monitored by the virtual router. This is usually an upstream interface. Should the interface go down, another router, configured as backup within the VRRP group, may take over as a master.
`<1-253>`	Priority delta value. The value by which the virtual router decrements its priority value during a circuit failover event. Configure this value to be greater than the difference of priorities between the master and backup routers.

Default

By default, circuit failover is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to configure circuit failover for the VRRP session with VR ID 1. Interface eth1 is considered the monitored interface.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#circuit-failover eth1 30
```

debug vrrp

Use this command to specify debugging options for VRRP.

Use the `no` parameter with this command to disable debugging.

Command Syntax

```
debug vrrp (all|event|packet [send|recv|])
no debug vrrp (all|event|packet [send|recv|])
```

Parameters

`all`	Specify debugging options for all VRRP events.
`event`	Specify debugging options for VRRP event troubleshooting.
`packet`	Specify debugging options for VRRP packets
`send`	Specify the debug option set for sent packets.
`recv`	Specify the debug option set for received packets.

Command Mode

Configure mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to enable all VRRP debug options.

```
#configure terminal
(config)#debug vrrp all
```

The example below shows how to enable debugging options for VRRP events.

```
#configure terminal
(config)#debug vrrp events
```

The example below shows how to enable debug options for VRRP packets sent.

```
#configure terminal
(config)#debug vrrp packet send
```

The example below shows how to enable debug options for VRRP packets received.

```
#configure terminal
(config)#debug vrrp packet recv
```

disable

Use this command to disable a VRRP session on the router (to stop the router from participating in virtual routing). When this command is configured, a backup Router assumes the Role of Master depending on its priority.

Command Syntax

```
disable
```

Parameters

None

Default

By default, VRRP session on the router is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to disable a VRRP session.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#disable
```

enable

Use this command to enable a VRRP session on the router (to make it participate in virtual routing). To make any changes to the VRRP configuration, first disable the router from participating in virtual routing using the `disable` command.

Note: Configure the virtual IP address and define an interface for the VRRP session (using the `virtual-ip` and `interface` commands) before using this command.

Command Syntax

```
enable
```

Parameters

None

Default

By default, VRRP session on the router is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to enable a VRRP session with VR ID 1 on interface eth0.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#enable
```

operational-ip

Use this command to set the primary IPv4 address.

Use the no parameter with this command to remove a primary IPv4 address.

Command Syntax

```
operational-ip A.B.C.D
no operational-ip
```

Parameters

A.B.C.D	IPv4 address.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#operational-ip 1.2.3.4
```

preempt-mode

Use this command to configure preempt mode. If set to `true`, the highest priority backup is always the master when the default master is unavailable. If set to `false`, a higher priority backup does not preempt a lower priority backup which is acting as master.

When the master router fails, the backup routers come online in priority order — highest to lowest. Preempt mode set to `true` allows a higher priority backup router to relieve a lower priority backup.

By default, a preemptive scheme is enabled whereby a higher priority backup virtual router that becomes available takes over for the backup virtual router that was elected to become master virtual router. This preemptive scheme can be disabled using the `preempt-mode false` command. If preemption is disabled, the backup virtual router that is currently elected as Master does not transition to backup again when the alternate backup router with higher priority becomes available.

Command Syntax

```
preempt-mode (true|false)
```

Parameters

`true`	Specify that preemption is enabled.
`false`	Specify that preemption is disabled.

Default

By default, preempt mode is true

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to enable the preempt mode.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#preempt-mode false

#configure terminal
(config)#router vrrp 1 eth0
(config-router)#preempt-mode true
```

priority

Use this command to configure the priority of a virtual router. The value of 255 (decimal) is reserved for the router that owns the IPvX address associated with the virtual router. The value of 0 (zero) is reserved for the Master router to indicate it is releasing responsibility for the virtual router. Higher values indicate higher priorities. Priority value to be used by this VRRP router in Master election

If the priority is high, the highest priority acts as a master. If the priority is low, lowest priority which is acting as backup.

Use the `no` parameter with this command to remove a priority and set the default value.

Command Syntax

```
priority <1-254>
no priority
```

Parameter

`<1-254>`	Priority of a virtual router.

Default

100 (non-owner)

255 (owner)

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

The example below shows how to configure priority 50 for the virtual router with VR ID 2 on interface eth0.

```
#configure terminal
(config)#router ip vrrp 2 eth0
(config-router)#priority 200
```

router vrrp

Use this command to enable a VRRP routing process.

Use the `no` form of this command to disable a VRRP routing process.

Command Syntax

```
router vrrp <1-255> IFNAME
no router vrrp <1-255> IFNAME
```

Parameters

`<1-255>`	Virtual router identifier
`IFNAME`	Interface name

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router vrrp 1 eth0
```

show debugging vrrp

Use this command to display the set VRRP debugging option.

Command Syntax

```
show debugging vrrp
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging vrrp
VVRRP debugging status:
VRRP event debugging is on
VRRP packet debugging is on
```

show running-config router vrrp

Use this command to show the running configuration for VRRP.

Command Syntax

```
show running-config router vrrp
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows the running configuration of VRRP. Virtual Router is configured as Master and Owner of IP address.

```
#show running-config router vrrp
!
router vrrp 1 eth0
 virtual-ip 39.0.0.24 owner
 advertisement-interval 5
 preempt-mode false
 enable
!
```

show vrrp

Use this command to display a list of virtual router identifiers that are configured on the router.

Command Syntax

```
show vrrp
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show vrrp

R1#show vrrp
VrId <1>
State is Master
Virtual IP is 10.10.12.6 (IP owner)
Interface is eth0
Priority is 255
Advertisement interval: 5 centi sec
Preempt mode is TRUE

R2#show vrrp
VrId <1>
State is Backup
Virtual IP is 10.10.12.6 (Not IP owner)
Interface is eth0
Priority is 100
Advertisement interval: 5 centi sec
Preempt mode is TRUE
```

Table 1-48 Explains the show command output fields.

Table 1-48: show vrrp output fields

Field	Description
VrID	Type of vr identifier by the system on the interface.
State	VRRP State: Master — The interface is acting as the master router interface. Backup —The interface is acting as the backup router interface.
Virtual IP	List of virtual IP addresses configured on the interface.
Interface	Name of the logical interface.

Field	Description
Priority	Configured VRRP priority for the interface.
Advertisement interval	Configured VRRP advertisement interval.
Preempt mode	Whether preemption is allowed on the interface.

show vrrp summary

Use this command to display a brief list of virtual router identifiers that are configured on the router.

Command Syntax

```
show vrrp summary
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS-SP version 4.2.

Example

```
#show vrrp summary
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

Interface  Admin  IpVersion  VRGroup  VRState  Priority    VIP          InterfaceIp
eth3       Up     IPv4      255      Master   255         66.0.0.1     66.0.0.1
eth3       Up     IPv6        1      Master   250         fe80::1      fe80::2
Total Number of IPv4 VRRP group 1
Total Number of IPv6 VRRP group 1
```

Table 1-49 explains the show command output fields.

Table 1-49: show vrrp summary output fields

Field	Description
interface	Interface name on which vrrp session is active.
Admin	Administrative state of the session Up/Down.
IpVersion	Internet protocol version IPv4 or IPv6.
VRGroup	Vrrp Group the router is part of.
priority	Configured VRRP priority for the interface.
VRState	VRRP State: Master - The interface is acting as the master router interface. Backup -The interface is acting as the backup router interface.
VIP	List of virtual IP addresses configured on the interface.
InterfaceIp	Primary IP address on the interface.

show vrrp <1-255>

Use this command to display VRRP information for a virtual router.

Command Syntax

```
show vrrp <1-255> IFNAME
```

Parameters

<1-255>	Virtual router identifier
IFNAME	Interface name

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
R1#show vrrp 7 xe1
VRRP Version: 3

VMAC enabled

Backward Compatibility disabled

Address family IPv4

VRRP Id: 7 on interface: xe1

State: AdminUp - Master

Virtual IP address: 10.10.10.81 (Owner)

Priority is 255

Advertisement interval: 100 centi sec

Master Advertisement interval: 100 centi sec

Skew time: 0 centi sec

Accept mode: FALSE

Preempt mode: TRUE

Multicast membership on IPv4 interface xe1: JOINED
```

```
R2#show vrrp 7 xe1

VRRP Version: 3

VMAC enabled

Backward Compatibility disabled

Address family IPv4

VRRP Id: 7 on interface: xe1

State: AdminUp - Backup

Virtual IP address: 10.10.10.81 (Not-Owner)

Priority is 100

Advertisement interval: 100 centi sec

Master Advertisement interval: 100 centi sec

Skew time: 0 centi sec

Accept mode: FALSE

Preempt mode: TRUE

Multicast membership on IPv4 interface xe1: JOINED
```

show vrrp (global | ipv4) statistics

Use this command to display VRRP global or ipv4 router statistics.

Command Syntax

```
show vrrp (global | ipv4 ) statistics
```

Parameters

`global`	Global (VRRP Router)
`ipv4`	VRRP IPv4 router

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.8.

Example

```
#show vrrp global statistics
 RRP Global Statistics
 Checksum Errors  : 0
 Version Errors   : 0
 VRid Errors      : 0
 Discontinuity Time: 00 hour, 00 min, 00 sec


#sh vrrp ipv4 statistics
Address family IPv4
VRRP Id: 1 on interface: xe1
Master Transitions: 0
Advertisements Rcvd: 0
Pkts Rcvd with IP TTL Errors: 0
Pkts Rcvd with Zero Priority: 0
Pkts Sent with Zero Priority: 0
Pkts Rcvd with Invalid TYPE: 0
Pkts Rcvd with Packet Length Errors: 0
Pkts Rcvd with IP Count Mismatch: 0
Discontinuity Time: 00 hour, 00 min, 00 sec
Refresh Rate: 1000 ms
```

snmp restart vrrp

Use this command to restart SNMP in Virtual Routing Redundancy Protocol (VRRP).

Command Syntax

```
snmp restart vrrp
```

Parameters

None

Default

By default, SNMP restart is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#snmp restart vrrp
```

switch-back-delay

Use this command to set a switch-back delay timer for the master VRRP router. This feature prevents the original master VRRP router from transitioning back to the master state after coming back online until the configured delay timer has expired.

Command Syntax

```
switch-back-delay <1-500000>
no switch-back-delay
```

Parameters

<1-500000>	Specify a switch-back delay in milliseconds.

Command Mode

Router mode

Default

By default, the switch-back delay is set to 0

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to set a switch-back delay timer of 7000 milliseconds.

```
#configure terminal
(config)#router vrrp 5 eth1
(config-router)#switch-back-delay 7000
```

virtual-ip

Use this command to set the Virtual Internet Protocol (VIP) address for the VRRP virtual router as VRRP Owner. This is the IP address used by end hosts to address their default gateway.

The VRRP Owner of the VIP address only responds to packets destined to the VIP address (for example, ICMP packets destined to the VIP address).

Use the `no` parameter with this command to remove a virtual IP address assignment.

Note: When using VRRP over MLAG, it is recommended to keep the VIP address different from the interface addresses. Having the same VMAC on both routers would essentially mean that they are both acting as masters for the VRRP group, which can lead to inconsistent routing behaviors.

Syntax Description

```
virtual-ip A.B.C.D (owner|)
no virtual-ip
```

Parameters

`A.B.C.D`	Specify the virtual IP address of the interface that participates in virtual routing.
`owner`	Specify the IP address as the owner.

Default

No default value is specified

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to configure the router as VRRP owner.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#virtual-ip 10.10.20.30 owner
```

The example below removes the virtual IP address assignment.

```
#configure terminal
(config)#router vrrp 1 eth0
(config-router)#no virtual-ip
```

vrrp compatible-v2

Use this command to enable the backward-compatibility feature. When enabled, both VRRPv3 and VRRPv2 inter-operation are supported.

Command Syntax

```
vrrp compatible-v2 (enable| disable)
```

Parameters

`enable`	Enable VRRPv2 inter-operation
`disable`	Disable VRRPv2 inter-operation

Default

By default, vrrp compatible-v2 is enabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#vrrp compatible-v2 enable
```

vrrp vmac

Use this command to enable or disable Virtual MAC (VMAC).

This command affects all VRRP groups in a router. On a single network segment, multiple VRRP groups can be configured, each using a different VMAC. The use of VMAC addressing allows for faster switchover when a backup router assumes the master role. When this command is used to enable a VMAC, the virtual router forwards packets with a special-purpose multicast VMAC address (0:0:5e:0:01:<VR ID>). Otherwise, it forwards with is interface's physical address.

The VMAC address is assigned to a router interface at the time the VRRP group is enabled in the router.

Command Syntax

```
vrrp vmac (enable|disable)
```

Parameters

`enable`	Enable virtual MAC addressing.
`disable`	Disable virtual MAC addressing and use physical MAC addressing.

Default

By default, VMAC is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to enable a virtual MAC address on the router.

```
#configure terminal
(config)#vrrp vmac enable
```

The example below shows how to disable a virtual MAC address on the router.

```
#configure terminal
(config)#vrrp vmac disable
```

CHAPTER 2 VRRP v6 Commands

This chapter describes the commands for VRRP IPv6.

- advertisement-interval
- circuit-failover
- disable
- enable
- preempt-mode
- priority
- router ipv6 vrrp
- router ipv6 vrrp vlan
- show running-config router ipv6 vrrp
- virtual-ipv6

advertisement-interval

Use this command to configure the advertisement interval of a virtual router. This is the length of time, in seconds, between each advertisement sent from the master to its backup(s). The master virtual router sends VRRP advertisements to other VRRP routers in the same group. The IPV6 VRRP advertisements are sent to the multicast address assigned to the VRRP IPV6 group (FF02:0:0:0:0:0:XXXX:XXXX) and a backup virtual router has to join all multicast groups within this range. As a convenient assignment, OcNOS sends a VRRP advertisement to the multicast address FF02::12.The advertisements are sent every second by default.

Note: VRRP Master router and backup routers should be configured with the same advertisement interval. If there is a mismatch in the configuration, VRRP goes to the INIT state.

Use the `no` parameter with this command to restore the default setting.

Command Syntax

```
advertisement-interval <5-4095>
no advertisement-interval
```

Parameter

`<5-4095>`	Specify the advertisement interval in centi-seconds (multiples of 5) when VRRPv3 is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to configure an advertisement interval of 10 seconds for the virtual router with VR ID 3 on interface eth0.

```
#configure terminal
(config)#router ipv6 vrrp 3 eth0
(config-router)#advertisement-interval 10
```

circuit-failover

Use this command to enable the VRRP circuit failover feature.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
circuit-failover IFNAME <1-253>
no circuit-failover IFNAME <1-253>
no circuit-failover (IFNAME|)
```

Parameters

`IFNAME`	Specify an interface of the router that is monitored by the virtual router. This is usually an upstream interface. Should the interface go down, another router, configured as backup within the VRRP group, may take over as a master.
`<1-253>`	Specify the delta value. The value by which the virtual router decrements its priority value during a circuit failover event. Configure this value to be greater than the difference of priorities between the master and backup routers.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows how to configure circuit failover for the VRRP session with VR ID 1. Interface eth1 is considered the monitored interface.

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#circuit-failover eth1 30
```

disable

Use this command to disable a VRRP session on the router (to stop the router from participating in virtual routing). Refer to enable to enable a VRRP session on the router.

When this command is configured, a backup router assumes the role of master depending on its priority.

Command Syntax

```
disable
```

Parameters

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows to disable VRRP session.

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#disable
```

enable

Use this command to enable a VRRP session on the router, so the router participates in virtual routing. To make any changes to the VRRP configuration, first disable the Router from participating in Virtual Routing using the `disable` command.

Note: Configure the virtual IP address and define an interface for the VRRP session (using the `virtual-ip` and `interface` commands) before using this command.

Command Syntax

```
enable
```

Parameters

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows to enable VRRP session with VR ID 1 on interface eth0

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#enable
```

preempt-mode

Use this command to configure preempt mode. If set to `true`, the highest priority backup is always the master when the default master is unavailable. If set to `false`, a higher priority backup does not preempt a lower priority backup that is acting as master.

When the master router fails, the backup routers come online in priority order — highest to lowest. Preempt mode set to `true` allows a higher priority backup router to relieve a lower priority backup.

By default, a preemptive scheme is enabled whereby a higher priority backup virtual router that becomes available takes over for the backup virtual router that was elected to become master virtual router. This preemptive scheme can be disabled using the `preempt-mode false` command. If preemption is disabled, the backup virtual router that is currently elected as Master does not transition to backup again when the alternate backup router with higher priority becomes available.

Command Syntax

```
preempt-mode true
preempt-mode false
```

Parameters

`true`	Specify that preemption is enabled.
`false`	Specify that preemption is disabled.

Default

Default is `true`.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows to configure preempt mode as False.

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#preempt-mode false
```

priority

Use this command to configure the VRRP router priority within the virtual router. Priority determines the role that each VRRP router plays and what happens if the master virtual router fails. If a VRRP router owns the IP address of the virtual router and the IP address of the physical interface, this router functions as the master virtual router.

Priority also determines whether a VRRP router functions as a backup virtual router and the order of ascendancy to becoming a master virtual router if the master virtual router fails.

Use the `no` parameter with this command to disable this feature.

Command Syntax

```
priority <1-255>
no priority
```

Parameter

`<1-255>`	Specify a priority. For the master router, specify 255; otherwise use any number in the range <1-254>.

Default

Default values for priority are:

- master router = 255
- backup = 100

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The example below shows to set the priority

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#priority 101
```

router ipv6 vrrp

Use this command to associate an IPv6 interface with a VRRP session. When issued, this command enters the `Router` mode.

Use the `no` parameter with this command to remove the IPv6 VRRP configuration. Disable the IPv6 VRRP session before using this command.

Command Syntax

```
router ipv6 vrrp <1-255> IFNAME
no router ipv6 vrrp <1-255> IFNAME
```

Parameters

`<1-255>`	Specify the ID of the virtual router session to create.
`IFNAME`	Specify the name of the IPv6 interface on which VRRP is enabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to enable an IPv6 VRRP session with VR ID 1 on interface eth0.

```
#configure terminal
(config)#router ipv6 vrrp 1 eth0
(config-router)#
```

router ipv6 vrrp vlan

Use this command to associate a VLAN with a VRRP session. When issued, this command enters the `Router` mode.

Use the `no` parameter with this command to remove the IPv6 VRRP configuration. Disable the IPv6 VRRP session before using this command.

Command Syntax

```
router ipv6 vrrp <1-255> vlan <1-4094>
no router ipv6 vrrp <1-255> vlan <1-4094>
```

Parameters

`<1-255>`	Specify a virtual router identifier. Must be unique for each routing process.
`<1-4094>`	Specify the actual VLAN identifier

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows the use of the `router vrrp` command to enter router mode. Note the change in the prompt.

```
#configure terminal
(config)#router ipv6 vrrp 100 vlan 123
(config-router)#
```

show running-config router ipv6 vrrp

Use this command to show the running configuration for VRRP.

Command Syntax

```
show running-config router ipv6 vrrp
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows the running configuration of IPv6 VRRP. Virtual Router is configured as Master and Owner of IP address.

```
#show running-config router ipv6 vrrp
!
router ipv6 vrrp 1 eth0
 virtual-ip fe80::202:b3ff:fed5:983e master
 circuit-failover eth1 30
 advertisement-interval 6
 preempt-mode false
 enable
!
#
```

virtual-ipv6

Use this command to set the virtual IPv6 address for the VRRP virtual router. This is the IPv6 address used by end hosts to address their default gateway. A maximum of one link-local and one global IPv6 address can be configured as virtual IP.

The VRRP Owner of the Virtual IPv6 address only responds to packets destined to the Virtual IPv6 address.

Use the no parameter with this command to remove a virtual IPv6 address assignment.

Syntax Description

```
virtual-ipv6 X:X::X:X (owner|)
no virtual-ip (X:X::X:X |)
```

Parameters

`X:X::X:X`	Virtual IPv6 address of the interface that participates in virtual routing.
`owner`	Specify the IPv6 address as the owner.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The example below shows how to configure a router as VRRP Owner:

```
#configure terminal
(config)#router ipv6 vrrp 1 eth1
(config-router)#virtual-ipv6 fe80::1 owner
(config-router)#virtual-ipv6 2001::10
```

The example below removes the virtual IP address assignment.

```
#configure terminal
(config)#router ipv6 vrrp 1 eth1
(config-router)#no virtual-ipv6
```

Index

A

ABRs 1436
accept 736
accept-lifetime 961, 1562
adding multiple instances of same AS to administer route selection 287
advertisement interval, VRRP Commands
  advertisement-interval 1758
aggregate-address 1609
aggregate-nexthop-check 434
area
  authentication 1289
area authentication 1289
Area Border Router
  configuring 1087
Area Border Router in OSPFv3 1203
Area Border Routers 1436
area default cost 1290
area default-cost 1433
area filter-list 1295
area nssa 1296, 1434
area range 1298, 1436
area stub 1300, 1437
area virtual-link 1301, 1438
area-password 965
area-password command 965
authentication 711
authentication commands
  accept-lifetime 1562
  ip rip authentication key-chain 1573
  ip rip authentication mode md5 1574
  ip rip authentication string 1575
  key chain 1582
  send-lifetime 1594
authentication key-chain 966
authentication mode md5 967
authentication OSPF 1105
authentication rip multiple keys 1543
authentication rip-multiple keys 1550
authentication send-only 968
auto-cost reference bandwidth 1304
auto-cost reference-bandwidth 1440
auto-summary 433

B

backward compatibility 1687
begin modifier 33
bfd 738, 739
bfd all-interfaces 1305
BFD Configuration 672
  Echo Function 672
  Echo Interval 679
  Multi-hop Peer Timer 675
  Single-hop Session Timer 677
  Slow Timer 674
  Topology 672
bfd echo 740
bfd echo interval 741
bfd interval 743
bfd multihop-peer 744
bfd multihop-peer A.B.C.D interval 746
bfd multihop-peer X:X::X:X interval 749
bfd notification 750
BFD protocol commands 774
  debug ospf bfd 779
  IS-IS
    isis bfd 781
BFD Protocol Configurations 682
bfd slow-timer 752
BFD Static Configurations 707
  Topology 707
BGP 1654
BGP Commands
  address-family 430
  aggregate-address 432
  auto-summary 433
  bgp aggregate-nexthop-check 434
  bgp always-compare-med 435
  bgp bestpath as-path ignore 437
  bgp bestpath compare-confed-aspath 439
  bgp bestpath compare-routerid 440
  bgp bestpath med 442
  bgp client-to-client reflection 445
  bgp cluster-id 446
  bgp confederation identifier 447
  bgp confederation peer 448
  bgp dampening 449
  bgp default local-preference 451
  bgp deterministic med 452
  bgp enforce-first-as 453
  bgp extended-asn-cap 454
  bgp fast-external-failover 455
  bgp g-shut 458
  bgp g-shut-capable 459
  bgp g-shut-local-preference 460
  bgp log-neighbor-changes 461
  bgp nexthop delay 463
  bgp nexthop enable 464
  bgp rfc1771-path-select 465
  bgp rfc1771-strict 466
  bgp router-id 467
  bgp scan-time 468
  bgp update-delay 471
  clear bgp * 473
  clear bgp A.B.C.D| X:X::X:X 472
  clear bgp external 480
  clear bgp peer-group 483
  clear ip bgp A.B.C.D
    clear ip bgp A.B.C.D 487
  clear ip bgp ASN 475, 479
  clear ip bgp dampening 479
  clear ip bgp flap-statistics 482

clear ip bgp view 487
debug bgp 492
distance 497
exit address family mode 498
ip as-path access-list 499
ip community-list 500, 504
ip community-list expanded 502
ip community-list standard 503
ip extcommunity-list
  expanded 507
  extended 506
  standard 505, 508
neighbor
  activate 515
  advertisement-interval 516
  capability dynamic 522
  capability graceful-restart 522
  capability orf prefix-list 523
  capability route-refresh 525
  collide-established 525
  default-originate 526
  distribute-list 531
  dont-capability-negotiate 532
  ebgp-multihop 533
  enforce-multihop 534
  fall-over bfd 537
  filter-list 537
  maximum-prefix 543
  next-hop-self 544
  peer-group add 548
  port 551
  remote-as 555
  remove private-as 557
  restart-time 558
  route-map 559
  route-reflector-client 561
  send-community 563
  shutdown 565
  soft-reconfiguration 566
  strict-capability-match 567
  timers 568
  unsuppress-map 570
  update-source 571
  version 572
  weight 573
neighbor attribute-unchanged 520
neighbor disallow-infinite-holdtime 530
neighbor g-shut 538
neighbor g-shut-timer 539
neighbor passive 547
neighbor prefix-list 553
network 576
redistribute 579
router bgp 582
show
  ip bgp attribute-info 652
  ip bgp cidr-only 652
  ip bgp community 608
  ip bgp community-info 654
  ip bgp community-list 610
  ip bgp filter-list 617
  ip bgp view 645
  ip protocols bgp 660
  show bgp neighbors received prefix-filter 629
  show bgp neighbors received-routes 630
  show bgp neighbors routes 631
  show bgp paths 636
  show bgp prefix-list 637
  show bgp quote-regexp 638
  show bgp route-map 640
  show bgp summary 643
  show ip bgp inconsistent-as 619
  show ip extcommunity-list 659
  synchronization 585
  timers 586
BGP community value
  command syntax 31
BGP configuration 214
  adding multiple instances of same AS to administer route selection 287
  configuring BGP Confederation 234
  configuring BGP distance 313
  configuring BGP Extended Community attributes 297
  configuring BGP Four-Byte AS 292
  configuring BGP weight per peer basis 317
  configuring Next-hop Tracking 304
  configuring Next-hop Tracking delay timer 307
  enabling BGP- different autonomous systems 215
  enabling EBGP Multihop 247
  enabling peer groups 275
  removing Multi-Exit Disc attribute from update messages 290
  route reflector 219
  route-map 217
bgp dampening 449
bgp g-shut 458
bgp g-shut-capable 459
bgp g-shut-local-preference 460
bgp inbound-route-filter 590
BGP VPN Commands
  bgp inbound-route-filter 590
  import map 593
  ip vrf 594
  neighbor
    send-community 597
  neighbor allow-egbp-vpn 595
  neighbor as-override 596
  neighbor send-community 597
  neighbor soo 598
  rd (route distinguisher) 599
  route-target 600
  show ip vrf 662
  site-of-origin 598
Bidirectional Forwarding Commands 735
Border Gateway BC Protocol (BGP) 1654

Border Gateway Protocol Configuration 213
braces
  command syntax 30

C

capability opaque 1308
capability restart 1309, 1423
capability restart graceful 1442, 1520
circuit-failover 1736, 1759
cisco-metric-behavior 1564, 1610
clear bgp
  * 473
  A.B.C.D|X:X::X:X 472
  external 480
  peer-group 483
clear bgp ipv6
  A.B.C.D| X:X::X:X 472
  external 480
  peer-group 483
clear clns is-neighbors 972
clear clns neighbors 971
clear data from RIPng routing table 1611
clear ip bgp
  ASN 475, 479
  view 487
clear ip ospf process 1312
clear ip rip route 1565
clear ip route kernel 104
clear ipv6 ospf process 1443
clear ipv6 rip route 1611
clear isis interface counter 976
clear isis process 977
clear isis route 973
collide-established 525
command abbreviations 29
command completion 29
command line
  errors 29
  help 28
  keyboard operations 32
command modes 36
  configure 36
  exec 36
  interface 36
  privileged exec 36
  router 36
command negation 30
command syntax
  ? 31
  . 31
  () 30
  {} 30
  | 30
  A.B.C.D/M 31
  AA:NN 31
  BGP community value 31
  braces 30
  conventions 30
  curly brackets 30
  HH:MM:SS 31
  IFNAME 31
  interface name 31
  IPv4 address 31
  IPv6 address 31
  LINE 31
  lowercase 30
  MAC address 31
  monospaced font 30
  numeric range 31
  parantheses 30
  parentheses 30
  period 31
  question mark 31
  square brackets 31
  time 31
  uppercase 30
  variable placeholders 31
  vertical bars 30
  WORD 31
  X:X::X:X 31
  X:X::X:X/M 31
  XX:XX:XX:XX:XX:XX 31
common commands 135
  match as-path 139
  match community 140
  match interface 142
  match ip address 143
  match ip address prefix-list 144
  match ip next-hop 145
  match ip next-hop prefix-list 146
  match ipv6 address 148
  match ipv6 address prefix-list 149
  match ipv6 next-hop 150
  match metric 153
  match origin 154
  match route-type 155
  match tag 156
  route-map 157
  set as-path 159
  set atomic-aggregate 160
  set comm-list delete 161
  set community 162
  set dampening 164
  set extcommunity 165
  set ip next-hop 168
  set ipv6 next-hop 169
  set level 170
  set metric 172
  set metric-type 173
  set origin 174
  set originator-id 175
  set tag 176
  set weight 177
compatible rfc1583 1313
Configuration
  BFD Configuration in BGP 698
  BFD Configuration in IS-IS 695

OSPF—BFD Multi-hop Session 691
configure
Area Border Router for OSPFv3 1203
BGP 214
configuring overload bit 856
distance 832
graceful restart 832
IS-IS 790
IS-IS-IPv6 861
metric 803
OSPF 1078
OSPFv3 1196
passive interface 856
RIPng 1558
route leaking 822
route summarization 827
route-map 217
configure mode 36
Configure Multipath eBGP 326
Configure Multipath iBPG 338
configure static RIPng routes 1626
configuring BGP Confederation 234
configuring BGP Extended Community attributes 297
configuring BGP Four-Byte AS 292
configuring OSPF
redistributing routes into OSPF 1091
configuring overload bit in IS-IS 856
configuring preempt mode 1741, 1762
configuring RIPng
enable RIPng 1558
Configuring VRRP
one virtual router 1668
two backup routers 1675
two virtual routers 1671
configuring VRRP IPv6 1723
cost
OSPF 1093
curly brackets
command syntax 30

D

damped route 650, 652
debug bfd all 753
debug bgp events 492
debug bgp filters 492
debug bgp fsm 492
debug bgp keepalives 492
debug bgp updates 492
debug ip ospf graceful-restart 1318, 1426
debug ipv6 ospf 1444
ifsm 1447
packet 1452
debug ipv6 ospf lsa 1448, 1449
debug ipv6 ospf nfsm 1450
debug ipv6 ospf nsm 1451
debug ipv6 ospf packet 1452
debug ipv6 ospf route 1455
debug ipv6 rip 1612
debug isis 978
debug ospf 1314
packet 1325
debug ospf database-timer rate-limit 1315
debug ospf events 1316
debug ospf ifsm 1317
debug ospf lsa 1322
debug ospf nfsm 1323
debug ospf nsm 1324
debug ospf packet 1325
debug ospf route 1327
debug ospf6 packet 1452
debug rib 108
debug vrrp 1737
default-information originate 980, 1328, 1456, 1569, 1613
default-metric 1330, 1458, 1570, 1614
default-metric command 1458
description 110
disable VRRP session 1738, 1760
distance 497, 1331, 1459, 1571, 1615
distance (OSPF command) 1331
distance in IS-IS 832
distribute-list 1267, 1460, 1572, 1616
domain-password 983

E

enable db-summary-opt 1334, 1462
enable IS-IS 790
enable multiple OSPF instances 1109
enable multiple OSPFv3 instances 1232
enable OSPFv3 1196
enable RIPng 1619
enable VRRP sessions 1739, 1761
exec command mode 36
exit-address-family 1468
exit-address-family command 498

F

fast-external-failover 455
fib retain 111

G

graceful restart 1048
graceful restart commands
neighbor capability graceful-restart 522
graceful restart in IS-IS 832

H

history 650, 652
host area 1338
how to disable VRRP session 1738, 1760
how to enable
authentication on an area 1105
authentication on an interface 1105
IS-IS 790

multiple OSPF instances 1109
multiple OSPFv3 instances 1232
OSPF on an interface 1078
OSPFv3 1196
RIP 1531
how to enable VRRP session 1739, 1761
how to enable/disable Virtual MAC 1756
how to enter router mode 1764
how to redistribute routes 1209
how to redistribute routes into OSPF 1091
how to set priority in OSPF 1082
how to set priority in OSPFv3 1199
how to set priority-IS-IS 793
how to set the Virtual IP address 1754
how to set the Virtual IPv6 address 1767
how to specify RIP version 1535

I

IFNAME 31
ignore-lsp-errors 989
import map 593
interface mode 36
Interior Gateway Protocol (IGP) 1654
Intermediate System to Intermediate System (IS-IS) 1651
ip bfd static all-interfaces 783
ip extcommunity-list
expanded 507
extended 506
standard 508
ip extcommunity-list standard 505
ip ospf authentication 1339
ip ospf authentication-key 1340
ip ospf bfd 1342
ip ospf cost 1343
ip ospf database-filter 1344
ip ospf dead-interval 1345
ip ospf disable all 1346, 1348
ip ospf hello-interval 1349
ip ospf message-digest-key 1351
ip ospf mtu 1353
ip ospf mtu-ignore 1354
ip ospf network 1355
ip ospf priority 1356
ip ospf retransmit-interval 1357
ip ospf transmit-delay 1358
ip rip authentication key-chain 1573
ip rip authentication mode md5 1574
ip rip authentication string 1575
ip rip receive version 1577
ip rip receive-packet 1576
ip rip send version 1579
ip rip send-packet 1578
ip rip split-horizon 1580
ip route 112
ip router isis 990
ip static bfd 785
ip static fall-over bfd 784
ip vrf 116, 594
IPv4 address
command syntax 31
IPv6 address
command syntax 31
ipv6 bfd static all-interfaces 786
IPv6 configuration 1723
ipv6 ospf cost 1468
ipv6 ospf dead-interval 1469
ipv6 ospf display route single-line 1470
ipv6 ospf link-lsa-suppression 1472
ipv6 ospf mtu-ignore 1473
ipv6 ospf neighbor 1474
ipv6 ospf network 1476
ipv6 ospf priority 1477
ipv6 ospf restart grace-period 1478, 1521
ipv6 ospf restart helper 1479, 1522
ipv6 ospf retransmit-interval 1481
ipv6 ospf6 transmit-delay 1482
ipv6 rip split-horizon 1618
ipv6 router mode 1764
ipv6 router ospf 1483
ipv6 router rip 1619
ipv6 static bfd 788
ipv6 static fall-over bfd 787
ipv6 te-metric 1485
isis authentication key-chain 993
isis authentication mode md5 994
isis authentication send-only 995
isis circuit-type 997
IS-IS commands 977
accept-lifetime 961
area-password 965
authentication key-chain 966
authentication mode md5 967
authentication send-only 968
clear clns is-neighbors 972
clear clns neighbors 971
clear isis interface counter 976
clear isis route 973
default-information originate 980
distance (IPv4) 982
isis authentication key-chain 993
isis authentication mode md5 994
isis authentication send-only 995
isis csnp-interval 998
isis hello-interval 1000
isis hello-multiplier 1001
isis metric 1005
isis network 1006
isis password 1007
isis priority 1008
isis restart grace-period 1044
isis restart helper 1046
isis restart-hello-interval 1045
isis retransmit-interval 1009
key 1017, 1581
key chain 1016
key-string 1017, 1018

max-area-address 1023
metric-style 1025
redistribute 1030
redistribute isis 1032
restart isis graceful 1048
send-lifetime 1035, 1038
set-overload-bit 1038
show clns is-neighbors 1051
show clns neighbors 1053
show cspf lsp 1055
show debugging isis 1055
show ip protocols 1061
show isis counter 1065
show isis counter command 1065
show isis database 1066
show isis interface 1069
show isis topology 1073
show running-config interface isis 1075
show running-config router isis 1076
summary-address 1040
ISIS Configuration 1651
IS-IS configuration 790
configuring distance 832
configuring metric 803
configuring overload bit 856
configuring passive interface 856
enabling IS-IS on an interface 790
graceful restart 832
L1 L2 area routing with multiple instances 816
L1 L2 area routing with single instance 810
redistributing routes into IS-IS 797
route leaking 822
route summarization 827
setting priority 793
IS-IS configuration-IPv6 861
isis csnp-interval 998
isis csnp-interval command 998
isis hello-interval 1000
isis hello-interval command 1000
isis hello-multiplier 1001
isis hello-multiplier command 1001
isis metric 1005
isis metric command 1005
isis network command 1006
isis password 1007
isis password command 1007
isis priority 1008
isis priority command 1008
isis restart grace-period 1044
isis restart helper 1046
isis restart-hello-interval 1045
isis retransmit-interval 1009
isis retransmit-interval command 1009
is-type 1015

K

key chain 755, 1016, 1582
key command 1017, 1581
key 754
key-string 1017, 1018, 1583

L

L1L2 area routing 810, 816
LINE 31
log-neighbor-changes command 461
lsp-gen-interval 1020
lsp-refresh-interval 1022

M

MAC address
command syntax 31
Match and Set Commands
match as-path 139
match community 140
match interface 142
match ip address 143
match ip address prefix-list 144
match ip next-hop 145
match ip next-hop prefix-list 146
match ipv6 address 148
match ipv6 address prefix-list 149
match ipv6 next-hop 150
match metric 153
match origin 154
match route-type 155
match tag 156
set as-path 159
set atomic-aggregate 160
set comm-list delete 161
set community 162
set dampening 164
set extcommunity 165
set ip next-hop 168
set ipv6 next-hop 169
set level 170
set metric 172
set metric-type 173
set origin 174
set originator-id 175
set tag 176
set weight 177
match as-path 139
match command
origin 154
match community 140
match interface 142
match ip address 143
match ip address prefix-list 144
match ip next-hop 145
match ip next-hop prefix-list 146
match ipv6 address 148
match ipv6 address prefix-list 149
match ipv6 next-hop 150
match metric 153
match origin 154

match route-type 155
match tag 156
max-area-address 1023
max-concurrent-dd 1360, 1487
maximum-area 1361
maximum-paths 121
maximum-prefix 1584
max-lsp-lifetime 1024
max-static-routes 122
MED 435
metric in IS-IS 803
metric-style command 1025
Multi Exit Discriminator 435
Multicast Commands
  show ipv6 rpf 127

N

neighbor 1364, 1585
  passive 547
  peer-group add 548
  remove-private-AS 557
  unsuppress-map 570
neighbor allow-egbp-vpn 595
neighbor as-override 596
neighbor attribute-unchanged 520
neighbor command 1620
  advertisement-interval 516
  capability dynamic 522
  capability graceful-restart 522
  capability orf prefix-list 523
  enforce-multihop 534
neighbor disallow-infinite-holdtime 530
neighbor g-shut 538
neighbor g-shut-timer 539
neighbor send-community 597
neighbor soo 598
net 1027
network 1586
network area 1365
network command 576
NSM Commands
  show router-id 129

O

offset-list 1587, 1621
Open Shortest Path First (OSPF) 1644
origin codes 650, 653
ospf abr-type 1367
OSPF commands
  area authentication 1289
  area default-cost 1290
  area filter-list 1295
  area nssa 1434
  area range 1298
  area stub 1300
  area virtual-link 1301
  auto-cost reference-bandwidth 1304
  bfd all-interfaces 1305
  capability opaque 1308
  capability restart 1309, 1423
  clear ip ospf process 1312
  compatible rfc1583 1313
  debug ospf 1314
  debug ospf database-timer rate-limit 1315
  debug ospf events 1316
  debug ospf ifsm 1317
  debug ospf ism 1317
  debug ospf lsa 1322
  debug ospf nfsm 1323
  debug ospf nsm 1324
  debug ospf packet 1325
  debug ospf route 1327
  default-information originate 1328, 1456
  default-metric 1330, 1458
  distance 1331, 1459
  distribute-list 1332
  enable db-summary-opt 1334
  host area 1338
  ip ospf authentication 1339
  ip ospf authentication-key 1340
  ip ospf bfd 1342
  ip ospf cost 1343
  ip ospf database-filter 1344
  ip ospf dead-interval 1345
  ip ospf disable all 1346, 1348
  ip ospf hello-interval 1349
  ip ospf message-digest-key 1351
  ip ospf mtu 1353
  ip ospf mtu-ignore 1354
  ip ospf network 1355
  ip ospf priority 1356
  ip ospf transmit-delay 1358
  ipv6 ospf mtu-ignore 1473
  max-concurrent-dd 1360
  maximum-area 1361
  neighbor 1364
  network area 1365
  ospf abr-type 1367
  ospf router-id 1372
  overflow database 1373
  overflow database external 1374
  passive-interface 1375, 1488
  redistribute 1376
  restart helper 1371, 1428
  restart ospf graceful 1378, 1429
  router ospf 1379
  show debugging ospf 1380
  show ip ospf 1381
  show ip ospf brder-routers 1385
  show ip ospf igp-shortcut-lsp 1395
  show ip ospf igp-shortcut-route 1396
  show ip ospf interface 1397
  show ip ospf neighbor 1402
  show ip ospf route 1406
  show ip ospf virtual-links 1409
  show ip protocols ospf 1411

summary-address 1417
timers lsa arrival 1418
timers throttle lsa 1420
OSPF Configuration 1644
OSPF configuration 1078
Area Border Router 1087
configuring multiple OSPF instances on same subnet 1120
configuring OSPF as PE-CE protocol for VPNs 320
configuring virtual links 1101
enable multiple OSPF instances 1109
enabling authentication 1105
enabling OSPF on an interface 1078
OSPF cost 1093
redistributing routes into OSPF 1091
setting priority 1082
ospf router-id 1372
OSPF VPN Commands
router ospf vrf 1527
OSPFv3 commands
abr-type 1432
area default-cost 1433
area nssa 1434
area range 1436
area stub 1437
area virtual-link 1438
auto-cost reference bandwidth 1440
capability restart graceful 1442, 1520
clear ipv6 ospf process 1443
debug ipv6 ospf 1444
debug ipv6 ospf ifsm 1447
debug ipv6 ospf lsa 1448, 1449
debug ipv6 ospf nfsm 1450
debug ipv6 ospf nsm 1451
debug ipv6 ospf packet 1452
debug ipv6 ospf route 1455
default-metric 1458
distribute-list 1460
enable db-summary-opt 1462
exit-address-family 1468
ipv6 ospf cost 1468
ipv6 ospf dead-interval 1469
ipv6 ospf display route single-line 1470
ipv6 ospf link-lsa-suppression 1472
ipv6 ospf neighbor 1474
ipv6 ospf network 1476
ipv6 ospf priority 1477
ipv6 ospf restart grace-period 1478, 1521
ipv6 ospf restart helper 1479, 1522
ipv6 ospf retransmit-interval 1481
ipv6 ospf transmit-delay 1482
ipv6 ospf6 transmit-delay 1482
ipv6 router ospf 1483
ipv6 te-metric 1485
max-concurrent-dd 1487
passive-interface 1488
redistribute 1489
restart ipv6 ospf graceful 1491, 1524
router ipv6 ospf 1493
router-id 1492
show debugging ipv6 ospf 1494
show ipv6 ospf database 1497
show ipv6 ospf interface 1501
show ipv6 ospf neighbor 1503
show ipv6 ospf route 1506
show ipv6 ospf topology 1509
show ipv6 ospf virtual-links 1511
show ipv6 ospf6 interface 1501
show ipv6 vrf 1513
summary-address 1516
OSPFv3 configuration 1196
configuring Area Border Router 1203
cost 1215
enable multiple OSPFv3 instances 1232
enable OSPFv3 on an interface 1196
redistribute routes into OSPFv3 1209
setting priority 1199
virtual links 1225
overflow database 1373
overflow database external 1374
Overload Bit 838

P

parantheses
command syntax 30
parentheses
command syntax 30
passive interface in IS-IS 856
passive-interface 1375, 1488, 1588, 1622
period
command syntax 31
preempt-mode 1741, 1762
priority 1763
priority-delta 1680
privileged exec mode 36

Q

question mark
command syntax 31

R

rd (route distinguisher) 599
recv-buffer-size 1589, 1623
redistribute 1030, 1489, 1590, 1624
redistribute command 1376
redistribute isis 1032
redistribute routes 1091
redistribute routes in OSPFv3 1209
redistributing routes into IS-IS 797
redistributing routes into OSPF 1091
redundancy using VRRP and OSPF 1690
removing Multi-Exit Disc attribute from update messages 290
restart helper 1371, 1428
restart ipv6 ospf graceful 1491, 1524

restart isis 1048
restart isis graceful 1048
restart ospf graceful 1378, 1429
restart-time 558
rfc1771-path-select 465
rfc1771-strict 466
RIP Commands
  accept-lifetime 1562
  cisco-metric-behavior 1564
  clear ip rip route 1565
  debug rip 1568
  default-information originate 1569
  default-metric 1570
  distance 1571
  distribute-list 1572
  ip rip authentication key-chain 1573
  ip rip authentication mode md5 1574
  ip rip authentication string 1575
  ip rip receive packet 1576
  ip rip receive version 1577
  ip rip send packet 1578
  ip rip send version 1579
  ip rip send-packet 1578
  ip rip split-horizon 1580
  key chain 1582
  key-string 1583
  maximum-prefix 1584
  neighbor 1585
  network 1586
  offset-list 1587
  passive-interface 1588
  recv-buffer-size 1589
  redistribute 1590, 1624
  route 1592
  router rip 1593
  send-lifetime 1594
  show debugging rip 1596
  show ip protocols rip 1597
  show ip rip 1599
  show ip rip interface 1601
  timers 1606
  version 1607
RIP configuration
  enabling RIP 1531
  RIPv2 md5 authentication 1550
  RIPv2 text authentication-multiple keys 1543
  specifying the RIP version 1535
RIP PE-CE Commands
  show ip rip database vrf 1636
RIP version 1535
RIP VPN commands
  show ip rip database vrf 1636
RIPng commands
  aggregate-address 1609
  cisco-metric-behavior 1610
  clear ipv6 rip route 1611
  debug ipv6 rip 1612
  default-information-originate 1613
  default-metric 1614
  distance 1615
  distribute-list 1616
  ipv6 rip split-horizon 1618
  ipv6 router rip 1619
  neighbor 1620
  offset-list 1621
  passive-interface 1622
  recv-buffer-size 1623
  route 1625
  route-map 1626
  router ivp6 rip 1627
  show debugging ipv6 rip 1628
  show ipv6 protocols rip 1629
  show ipv6 rip interface 1631
  show ipv6 rip show ipv6 rip 1630
  timers basic 1632
RIPng configuration 1558
route 1592, 1625
route leaking in IS-IS 822
route summarization in IS-IS 827
route-map 157, 1626
router ipv6 ospf 1493
router ipv6 rip 1627
router isis 1034
router mode 36
router ospf 1379
router ospf vrf 1527
router rip 1593
router vrrp 1743
route-reflector 219
router-id 1492
route-target 600
Routing Information Protocol (RIP) 1646

S

scan-time 468
send-lifetime 1035, 1038, 1594
send-lifetime 758
set administrative distance 1615
set as-path 159
set atomic-aggregate 160
set comm-list delete 161
set community 162
set dampening 164
set extcommunity 165
set interface null0 167
set ip next-hop 168
set ipv6 next-hop 169
set level 170
set metric 172
set metric-type 173
set origin 174
set originator-id 175
set tag 176
set weight 177
set-overload-bit 1038
setting priority 1199
setting router priority 1763

setting-priority 793
show bfd 759
show bfd interface 760
show bfd session 762
show bfd session A.B.C.D 766
show bfd session ipv6 769
show bgp ipv6 620
show bgp neighbors received prefix-filter 629
show bgp neighbors received-routes 630
show bgp neighbors routes 631
show bgp paths command 636
show bgp prefix-list command 637
show bgp quote-regexp command 638
show bgp route-map command 640
show clns is-neighbors 1051
show clns is-neighbors command 1051
show clns neighbors command 1053
show commands 33
  exclude modifier 34
  include modifier 34
  redirect modifier 35
show cspf lsp 1055
show debugging bfd 772
show debugging ipv6 ospf 1494
show debugging ipv6 rip 1628
show debugging isis 1055
show debugging ospf 1380
show debugging rib 124
show debugging rip 1596
show debugging vrrp 1744
show ip bgp rtfilter all 657
show ip extcommunity-list 659
show ip ospf 1381
  border routers 1385
show ip ospf border-routers 1385
show ip ospf igp-shortcut-lsp 1395
show ip ospf igp-shortcut-route 1396
show ip ospf interface 1397
show ip ospf neighbor 1402
show ip ospf route 1406
show ip ospf virtual-links 1409
show ip protocols 1061
show ip protocols bgp 660
show ip protocols ospf 1411
show ip protocols rip 1597
show ip rip command 1599
show ip rip database vrf 1636
show ip rip interface 1601
show ipv6 isis route 1065
show ipv6 ospf database 1497
show ipv6 ospf database, router 1491, 1524
show ipv6 ospf interface 1501
show ipv6 ospf neighbor 1503
show ipv6 ospf route 1506
show ipv6 ospf topology 1509
show ipv6 ospf virtual-links 1511
show ipv6 ospf6
  interface 1501
show ipv6 rip interface 1631
show ipv6 rpf 127
show ipv6 vrf 1513
show isis counter 1065
show isis database command 1066
show isis interface 1069
show isis topology 1073
show router-id 129
show running-config as-path access-list 663
show running-config community-list 664
show running-config interface isis 1075
show running-config route-map 179
show running-config router isis 1076
show running-config router vrrp 1745
show running-config router-id 131
show running-config vrf 133
snmp restart bfd 773
soft-reconfiguration 566
square brackets
  command syntax 31
stale route 650, 652
summary-address 1040, 1417, 1516
suppressed route 650, 652
synchronization command 585

T

time
  command syntax 31
timers 1606
timers basic 1632
timers lsa arrival 1418
timers throttle lsa 1420
transmit delay 1397
transmit-delay 1482

U

UDLD configuration 83
update-delay 471

V

valid route 650, 653
version 1607
vertical bars
  command syntax 30
virtual IP address 1754
virtual IPv6 address 1767
virtual links 1101
Virtual MAC command 1756
Virtual routing and forwarding (VRF) 1642
virtual-ip 1754
virtual-ip command 1767
VRF 1527
VRF Configuration 1642
VRRP Commands 1732, 1757
  circuit-failover 1759
  debug vrrp 1737
  disable 1738, 1760

enable 1739, 1761
preempt-mode 1741, 1762
priority 1763
show debugging vrrp 1744
show running-config router vrrp 1745
virtual-ip 1754
vrrp vmac 1756
VRRP process 1666
VRRPv6 Commands
router ipv6 vrrp 1764
virtual-ip 1767

W

WORD 31